

Assessment of the Impact of Cosmic-Ray-Induced Neutrons on Hardware in the Roadrunner Supercomputer

Sarah E. Michalak, Andrew J. DuBois, Curtis B. Storlie, Heather M. Quinn, William N. Rust,
David H. DuBois, David G. Modl, Andrea Manuzzato, and Sean P. Blanchard

Abstract—Microprocessor-based systems are a common design for high-performance computing (HPC) platforms. In these systems, several thousands of microprocessors can participate in a single calculation that may take weeks or months to complete. When used in this manner, a fault in any of the microprocessors could cause the computation to crash or cause silent data corruption (SDC), i.e., computationally incorrect results that originate from an undetected fault. In recent years, neutron-induced effects in HPC hardware have been observed, and researchers have started to study how neutrons impact microprocessor-based computations. This paper presents results from an accelerated neutron-beam test focusing on two microprocessors used in Roadrunner, which is the first Petaflop supercomputer. Research questions of interest include whether the application running affects neutron susceptibility and whether different replicates of the hardware under test have different susceptibilities to neutrons. Estimated failures in time for crashes and for SDC are presented for the hardware under test, for the Triblade servers used for computation in Roadrunner, and for Roadrunner.

Index Terms—Failures in time (FIT), neutron-beam testing, silent data corruption (SDC), single-event effect, soft error.

I. INTRODUCTION

LARGE-SCALE scientific computations are frequently performed on high-performance computing (HPC) platforms. These computations can use thousands of processors and run for weeks to months. Many HPC platforms use commercial off-the-shelf (COTS) microprocessors, as opposed to radiation-hardened devices, for such computation. Neutron-induced effects in COTS microprocessors include single-event upsets (SEUs) in the caches, register files, pipeline registers, and memory; single-event transients (SETs) in functional units;

and single-event functional interrupts (SEFIs) in control logic. The results of such neutron-induced effects can include failures (e.g., system and application crashes) and silent data corruption (SDC), which occurs when an undetected error causes the system to deliver computationally incorrect results. Neutrons have been implicated in crashes and SDC in different architectures [1]–[7]. While alpha particles can lead to similar issues [8], [9], this paper focuses on the effects of neutrons since it is concerned with the experience of systems located at Los Alamos National Laboratory, which is at high elevation (7200 ft) and hence experiences a higher neutron flux than that at sea level.

Because SEUs, which include single-bit upsets (SBUs) and multi-bit upsets (MBUs), are increasingly noticeable in terrestrial applications, COTS microprocessor designers and system designers often include some protection from SEUs. These protections include error-correcting codes (ECC), bit interleaving in caches, and parity checks. For computations run on HPC platforms, software-level protections, such as checkpoint/restart, are also implemented. In checkpoint/restart, the calculation's state is periodically saved to hard disk so that a calculation can be restarted from the previous state if necessary. A second method that may be implemented is algorithm-based fault tolerance (ABFT), but it relies on specialized knowledge by the programmer or *ad hoc* optimization of the code by hand, e.g., [10], [11]. ABFT usually decreases performance, which is paramount in HPC systems, and is typically not used in HPC systems.

While these protections are useful, the failures and SDC that can result from neutron induced are not completely suppressed. HPC platforms used for scientific computation are particularly sensitive to neutron-induced faults due to their large size and the availability requirements of the applications that run on them. An HPC platform can contain thousands of replicates of a particular microprocessor or other sensitive device, making these large platforms more likely to experience the effects of cosmic-ray-induced neutrons than smaller systems.

With a calculation using many processors, it is often the case that if a single microprocessor crashes the entire calculation will be stopped, the previous checkpoint loaded, and the calculation restarted. For these situations, the application runtime is increased each time the calculation is restarted from a checkpoint since all of the runtime between the checkpoint and the crash is lost. Furthermore, SDC can be difficult to detect. Therefore, in HPC platforms neutron-induced effects are

Manuscript received September 23, 2011; revised January 26, 2012; accepted February 27, 2012. This work was supported in part by the Los Alamos National Security, LLC under Contract DE-AC52-06NA25396 and in part by the U.S. Department of Energy.

S. E. Michalak, C. B. Storlie, and W. N. Rust are with the Statistical Sciences Group, Los Alamos National Laboratory, Los Alamos, NM 87544 USA (e-mail: michalak@lanl.gov; storlie@lanl.gov; wnr@lanl.gov).

A. J. DuBois, D. H. DuBois, D. G. Modl, and S. P. Blanchard are with the Systems Engineering and Integration Group, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: ajd@lanl.gov; dhd@lanl.gov; digem@lanl.gov; seanb@lanl.gov).

H. M. Quinn is with the Space Data Systems Group, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: hquinn@lanl.gov).

A. Manuzzato was with Università degli Studi di Padova, 35122 Padova, Italy (e-mail: andrea.manuzzato@iee.org).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TDMR.2012.2192736

82 of concern since (1) system crashes affect application runtimes
83 and (2) SDC in scientific applications may lead to incorrect
84 scientific conclusions.

85 This paper presents results from neutron-beam testing of
86 hardware identical to that used in Roadrunner [12], the first
87 Petaflop system [13]. The hardware was tested while running
88 different applications including some used for scientific re-
89 search. The structure of this paper is as follows. Section II
90 discusses related microprocessor studies. Section III presents
91 the test setup, with Section IV detailing the results. Section V
92 offers conclusions from this work. The statistical methods for
93 the data are described in [14], while this work augments and
94 complements [15] by presenting data pertaining to permanent
95 failures and estimated failures in time (failures in 10^9 hours
96 of operation or FIT) for failures (e.g., application and system
97 crashes) and for SDC for the microprocessors tested and the
98 other hardware in their beampaths, for the Triblade server
99 used for computation in the Roadrunner platform, and for the
100 Roadrunner platform itself.

101 II. RELATED WORK

102 There is more than a decade's worth of static test data on
103 microprocessors [16]–[20], with a number of recent publica-
104 tions addressing more modern microprocessors with reduced
105 feature sizes or multiple processing cores [4], [21]–[24]. While
106 static testing is often the basis for error rate calculations, it
107 can be difficult to translate the errors from static tests into
108 dynamic error rates that reflect the field experience of real-
109 world systems. Determining the overall effect of radiation on
110 microprocessors is not simple, as faults in a system can remain
111 dormant for several thousands of clock cycles before triggering
112 an error or can be masked. In addition, the operating system and
113 any software being used can create noise in the system, making
114 it difficult to determine the cause of system crashes.

115 Several studies performed dynamic testing similar to that
116 performed here. These include [4] (SPARC64 V microproces-
117 sor), [5] (IBM POWER6 and Intel Core2 5160 Xeon Wood-
118 crest microprocessors), [7] (Intel Core2 5160 Xeon Woodcrest
119 microprocessor), [6] (IBM POWER6 microprocessor), and [2]
120 (Intel Itanium processor). All of these studies except [5] explic-
121 itly report observing SDC or events that could lead to SDC.
122 Further, [4], [5], and [7] studied whether different applications
123 led to differing susceptibilities to neutrons. [5] did not establish
124 differing susceptibilities, and while point estimates of logic der-
125 ating factors provided in [4] suggested some differences for the
126 applications used, 95% confidence intervals largely coincided.
127 [7] found that while the mean times to first indication of failure
128 (MTFIF) for some of the pairs of applications studied had
129 high probability of being different, the idle condition did not
130 have the highest MTFIF. Finally, [16]–[19] performed proton
131 testing of Pentium and Celeron microprocessors while running
132 consistency checks and a workload simulator, with both failures
133 and SDC observed.

134 III. TEST SETUP AND EXPERIMENTAL PROTOCOL

135 Hardware from Roadrunner was tested at Los Alamos Na-
136 tional Laboratory's (LANL) Los Alamos Neutron Science
137 Center (LANSCE) Irradiation of Chips and Electronics (ICE)

House in October 2009 to investigate the neutron susceptibility
138 of the two microprocessors used in Roadrunner along with the
139 hardware in their respective beampaths. Both microprocessors,
140 the IBM PowerXCell 8i (Cell) and the AMD Opteron 2210 HE,
141 have been commercially available. The test setup included test-
142 ing multiple replicates of the Cell microprocessor and Opteron
143 microprocessor while running different applications, including
144 some used for scientific computation.

145 The Cell microprocessors and Opteron microprocessors were
146 operated in the neutron beam in their field configuration in
147 a Triblade blade server. A Triblade [12] includes one IBM
148 LS21 blade, two IBM QS22 blades, and an expansion blade to
149 manage data traffic. The LS21 blade has two dual-core Opteron
150 2210 HE microprocessors, and the QS22 blades (QS22a and
151 QS22b) each have two Cell microprocessors.

152 The Cell is a 65nm silicon-on-insulator (SOI) microproces-
153 sor with 1 PowerPC processor element (PPE) that controls 8
154 synergistic processor elements (SPE); [25, p. 5] provides a
155 diagram of the Cell architecture. The 3.2 GHz PPE includes
156 a PowerPC processor unit (PPU) that is based on the PowerPC
157 architecture, a parity-protected 32 KB L1 data cache, a parity-
158 protected 32 KB L1 instruction cache, and a 512 KB L2
159 cache with ECC on data and parity on directory tags (which is
160 recoverable using redundant directories). Each 3.2 GHz SPE in-
161 cludes a synergistic processor unit (SPU) and an ECC-protected
162 256 KB dedicated non-caching local store. The QS22 blade
163 that housed the Cells during the testing included 8 GB of ECC
164 double data rate 2 (DDR2) dynamic random access memory
165 (DRAM).

166 The Opteron 2210 HE is a 1.8 GHz 90nm SOI dual-core
167 microprocessor. See [26, p. 2] for a diagram of the Opteron
168 2210 EE microprocessor, which has a design similar to the
169 Opteron 2210 HE tested at LANSCE. Each Opteron core has
170 an ECC-protected 64 KB L1 data cache, a parity-protected
171 64 KB L1 instruction cache, and an ECC-protected 1MB L2
172 cache. The LS21 blade that housed the Opteron 2210 HE for
173 the testing included 16 GB of ECC DDR2 DRAM.

174 A. Test Setup

175 There are two aspects of the test setup: the hardware test
176 setup and the software test setup. Both aspects were specifically
177 designed to mimic how the devices under test operate in the
178 field as part of the Roadrunner platform.

179 1) *Hardware Test Setup:* Due to their extreme size, most
180 HPC platforms need an efficient power, cooling and network
181 design, as the entire system might span thousands of square
182 feet. To this end, most platforms are designed to be housed
183 in server racks, with each rack housing multiple chassis. In
184 a blade-based platform such as Roadrunner, each chassis will
185 house multiple blades. In Roadrunner, the rack provides physi-
186 cal structure; the chassis provides a common interface to power,
187 network, and cooling; and the LS21 and QS22 compute blades
188 provide the compute infrastructure.

189 The hardware tested included four Triblades and a
190 BladeCenter-H (BC-H Type 8852) [27] chassis, which is de-
191 signed to house up to three Triblades. The testing required
192 additional hardware for system control and for neutron fluence
193

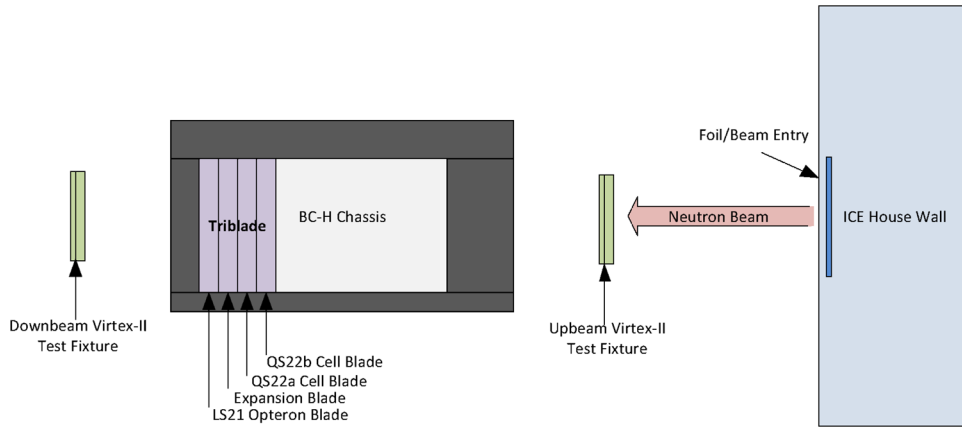


Fig. 1. Schematic of test setup with Triblade in the BC-H slot furthest from the beam source.

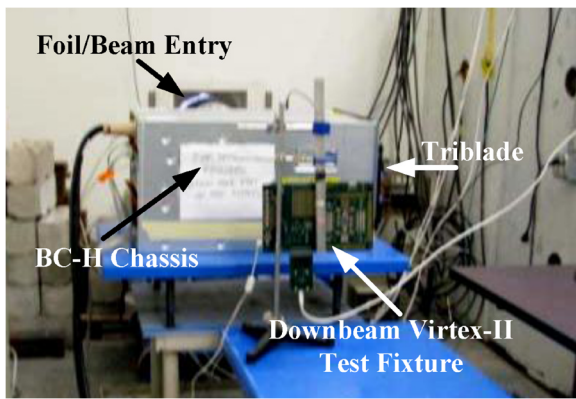


Fig. 2. Photo of test set-up.

194 exposure measurements. In many HPC platforms, a front-end
195 node is used to manage the back-end compute nodes or blades,
196 and likewise for this testing a front-end node was necessary to
197 control the system under test. Specifically, it was used to boot
198 the Triblades, start applications on the system under test and
199 to monitor its health, all of which were performed manually
200 by the experiment personnel. An IBM eServer X Series 336
201 provided this capacity. This server was placed in the user area
202 of the facility so that it would be protected from the neutron
203 beam.

204 Since the hardware setup included more physical matter
205 (chassis, metal enclosures, large heatsinks) than is typical, two
206 Xilinx Virtex-II [28] test fixtures [29] were included in the test
207 setup for calculating corrected neutron fluence exposures for
208 the hardware under test. Specifically, one was placed upbeam
209 of the BC-H and the other was placed downbeam of the BC-H.

210 Figs. 1 and 2 show the hardware test setup. Fig. 1 provides
211 a diagram of the test facility and the hardware under test,
212 including the point at which the beam enters the test facility,
213 the Virtex-IIs, the BC-H chassis and the Triblade under test.
214 Fig. 2 provides a complementary photograph with key aspects
215 of the test setup labeled. The BC-H was oriented so that with a
216 single Triblade under test, the beam first entered the QS22b,
217 followed by the QS22a, the expansion blade, and the LS21
218 respectively. For the testing, one Cell had a higher beam aim
219 than the other, with “Upper Cell” denoting this Cell and “Lower
220 Cell” denoting the Cell with the lower beam aim. For the two

Opertons in the LS21, “Upper Opteron ” and “Lower Opteron”
have analogous meanings.

221
222
223 2) *Software Test Setup*: The test applications for the Cell
224 included five computational test codes (hybrid Linpack [30], 224
[31], a correlator test code [32], a conjugate gradient solver,
225 VPIC [33], and an integer adder, which are further described
226 below) and an idle test code in which the Opteron interrogates
227 the Cell to determine if all processor elements (PPE and SPEs)
228 are all still operational. Hybrid Linpack performs the High-
229 Performance Linpack benchmark calculation, optimized for the
230 Triblade architecture with most of the computation performed
231 on the Cell. The correlator test code performs a multiply and
232 accumulate needed for certain radio-astronomy applications. It
233 utilizes both the Opteron and PPE in very limited ways, with
234 most of the computation performed on the SPEs. The conjugate
235 gradient method is a member of a family of iterative solvers
236 used primarily on large sparse linear systems arising from the
237 discretization of partial differential equations. The conjugate
238 gradient test code used here performs a double precision, pre-
239 conditioned conjugate gradient (CG) algorithm and utilizes the
240 Opteron primarily for generation of the sparse linear system,
241 with the CG implementation taking place on the Cell. VPIC is a
242 3-D electromagnetic relativistic particle-in-cell plasma physics
243 simulation code. The version used for this testing was written
244 to run on the Cell microprocessor in a hybrid microprocessor
245 environment like that of a Triblade. The integer add test code is
246 a simple hybrid code that executes primarily on the SPEs, using
247 vector integer units to perform simple adds. Vector registers
248 on the SPEs are loaded, vector adds are executed over these
249 registers and verified for correctness. A test with a Cell beam
250 aim used one of these applications or a test condition referred
251 to as “varied” in which the Cell executed two or more of the
252 applications.
253

The Opteron test applications included an Opteron-only ver-
sion of the correlator test code, idling, and running the Linux
top command, which is considered an idle condition in the
analyses that follow. The Opteron-only correlator test code
performs the multiply and accumulate described for the Cell
above on a single Opteron core, with both cores on the Opteron
under test running the code during the testing.

Each test code was designed so that it completed its work
in roughly one minute. The software setup was instrumented

263 to run the test code continuously and return output data each
 264 time the test code completed. The output data included start and
 265 stop times, the application being run, the hardware running it,
 266 whether an SDC occurred and with the Cell idle code whether
 267 the Cells under test were still responding.

268 It should be noted that initial testing of the devices to
 269 determine the static sensitivity of the caches and registers to
 270 SEU was not undertaken. All of the FIT estimates and other
 271 results determined from this study are based on the described
 272 dynamic system use.

273 B. Experimental Procedure

274 For a given experiment, a single Cell or Opteron was config-
 275 ured to run the desired application while the beam was aimed
 276 so that it irradiated all of the hardware in the Triblade and the
 277 BC-H in that microprocessor's beampath. With two QS22s in
 278 a Triblade, when a Cell in one QS22 is running an application,
 279 the corresponding Cell in the other QS22 is in the beampath.
 280 This second Cell in the beampath was set to run the Cell idle
 281 test code. Since the beam irradiated a cylindrical volume within
 282 the Triblade under test and the BC-H, certain attribution of an
 283 error to the Cells or Opteron in the beampath is not possible.
 284 In particular, other hardware in the beampath or hardware that
 285 was affected by scatter could be the cause of an observed error.
 286 Errors could also be the result of causes external to the beam,
 287 but this is much less likely.

288 The experimental protocol was to start the required test
 289 application on the appropriate microprocessor while the beam
 290 was off. Once the test application was observed to be operating
 291 properly (e.g., the test code had produced one or more output
 292 lines), the beam was started. The experiment continued until
 293 a state of system inoperability (e.g., a system or application
 294 crash) was reached or until sufficient time had elapsed. The
 295 beam was then turned off, data pertaining to neutron fluence
 296 exposure of the system under test were collected, and the sys-
 297 tem was rebooted before beginning the next test. For the Cells,
 298 the test procedure was to cycle through the test applications on a
 299 particular Cell, typically until it became inoperable. Repeating
 300 each test code periodically permits investigation of any aging or
 301 dose-related effects resulting from increasing exposure to the
 302 beam. The procedure for the Optérons, which received much
 303 less testing, was to use the Opteron-only correlator code and
 304 possibly an idle condition (idling or running the Linux top
 305 command). Functionality of the Opteron while it was idling or
 306 running the Linux top command was assessed by ascertaining
 307 its continued responsiveness.

308 In all, 113 experiments were performed, with 14 Cells and
 309 3 Optérons operated in the beam. The first three experiments,
 310 which were the only data collected for Triblade 2, were omitted
 311 from the results since these tests had three Tribblades in the
 312 beam whereas the remaining experiments had only a single
 313 Triblade in the beam. Another experiment with missing beam
 314 fluence data was also omitted from the analyzes. The Opteron
 315 beampath tests were performed after the Cell beampath tests
 316 since the Cells were of primary interest in the testing. Thus, the
 317 behavior of the Optérons and the hardware in their beampaths

TABLE I
 PROPORTION REDUCTION AT EACH OF FOUR BEAM AIMS

Upper Cell	Lower Cell	Upper Opteron	Lower Opteron
0.60	0.69	0.69	0.70
(0.56, 0.64)	(0.66, 0.72)	(0.56, 0.84)	(0.60, 0.79)

in a Triblade with no previous exposure to the beam cannot be
 estimated based on this testing.

Two different beam diameters were used for the experi-
 ments: a two-inch beam diameter for the first 53 experiments
 and a one-inch beam diameter for the remaining 59, where
 these beam diameter measurements reflect the full-width half-
 maximum (FWHM) boundary. All testing was performed at
 nominal voltages and nominal temperatures with the test fixture
 at normal incidence to the beam.

C. Corrected Neutron Fluences

Typically, corrected neutron fluences would be based on
 the decrease in flux given the distance from the beam source.
 Without the BC-H chassis and the Tribblades in the beam,
 the calculated decrease in flux from the beam source to the
 downbeam Virtex-II is about 20%.

The reduction in flux at each of the four beam aims de-
 scribed at the end of the Hardware Test Setup description in
 Section III.A (lower Cell, upper Cell, lower Opteron, upper
 Opteron) was estimated based on the experimental data to
 assess whether the BC-H and Triblade led to additional re-
 duction in beam intensity. Table I presents the posterior mean
 of the reduction in beam flux at each of the four beam aims
 and corresponding 95% credible intervals (CI). These values
 are based on the Virtex-II measurements from the upbeam
 and downbeam Virtex-II devices and a distance of 95 inches
 between the point at which the beam enters the test facility and
 the downbeam Virtex-II, which is the most common distance
 between these two points in the experimental data, and they
 incorporate both attenuation resulting from the material in the
 beam and divergence of the beam resulting from distance from
 the beam source. These results and those throughout this paper
 are based on the model described in Section IV-D, which
 permits different reductions in the beam at each beam aim and
 incorporates the uncertainty in these reductions; see [14] for
 details. The narrower CIs for the Cell beam aims reflect the
 greater number of experiments performed at the two Cell beam
 aims.

Since Table I demonstrates that the decrease in flux is larger
 than that expected due to only distance from the beam source,
 the neutron fluence exposures for different tests were corrected
 based on both distance from the beam source and attenuation
 through matter. The decrease in flux based on distance was
 calculated as usual, i.e., under the assumption that the beam
 is a point source with the reduction proportional to the squared
 distance from the beam source. The decrease in radiation due
 to attenuation through variable matter (i.e., the Triblade) is
 difficult, if not impossible, to account for precisely. However,
 as described in the previous paragraph the total proportion
 reduction through the entire Triblade and BC-H for each of the

TABLE II
HARDWARE NEUTRON EXPOSURE: AN ASTERISK (*) INDICATES A
QS22 THAT EXPERIENCED A VOLTAGE REGULATOR
FAILURE DURING THE TESTING

Blade	Beam Aim	Corrected Fluence ($\frac{neutrons}{cm^2}$)	
		Lower Bound	Upper Bound
Triblade 1 LS21	Upper Opteron	9.20×10^7	2.38×10^8
Triblade 1 LS21	Lower Opteron	3.42×10^8	8.94×10^8
Triblade 1 QS22a*	Upper Cell	1.28×10^9	2.56×10^9
Triblade 1 QS22a*	Lower Cell	7.98×10^8	2.08×10^9
Triblade 1 QS22b	Upper Cell	1.29×10^9	2.57×10^9
Triblade 1 QS22b	Lower Cell	8.01×10^8	2.09×10^9
Triblade 2 LS21	Upper Opteron	0	0
Triblade 2 LS21	Lower Opteron	0	0
Triblade 2 QS22a	Upper Cell	0	0
Triblade 2 QS22a	Lower Cell	8.16×10^8	2.13×10^9
Triblade 2 QS22b*	Upper Cell	0	0
Triblade 2 QS22b*	Lower Cell	8.18×10^8	2.13×10^9
Triblade 3 LS21	Upper Opteron	0	0
Triblade 3 LS21	Lower Opteron	0	0
Triblade 3 QS22a	Upper Cell	9.98×10^9	1.99×10^{10}
Triblade 3 QS22a	Lower Cell	1.86×10^9	4.85×10^9
Triblade 3 QS22b*	Upper Cell	1.00×10^{10}	1.99×10^{10}
Triblade 3 QS22b*	Lower Cell	1.86×10^9	4.86×10^9
Triblade 4 LS21	Upper Opteron	0	0
Triblade 4 LS21	Lower Opteron	4.37×10^8	1.14×10^9
Triblade 4 QS22a	Upper Cell	3.82×10^9	7.62×10^9
Triblade 4 QS22a	Lower Cell	1.43×10^{10}	3.73×10^{10}
Triblade 4 QS22b*	Upper Cell	3.83×10^9	7.64×10^9
Triblade 4 QS22b*	Lower Cell	1.43×10^{10}	3.74×10^{10}

four beam aims can be estimated with the Virtex-II readings. With this information, the neutron fluence to which a particular component is exposed prior to a particular error or operator decision to stop a test is assumed to lie between a lower bound and an upper bound explained below. The resulting uncertainty in neutron exposure of the component is explicitly incorporated in the model described in Section IV-D and reflected in the results presented here. The lower bound assumes that all of the attenuation caused by the beam passing through the BC-H and Triblade under test happened upstream of the component under test, while the upper bound assumes all attenuation happened downstream of the component under test. While neither of these cases may reflect the actual reduction due to attenuation of the beam, they best capture the knowledge of beam attenuation that resides in the experimental data without making any further assumptions. See [14] for details of the model used.

Table II details the posterior means of the upper bounds and lower bounds of the corrected neutron fluence, for neutrons with energies greater than 10 MeV, accumulated at each beam aim during the testing. The corrected fluences are based on the posterior mean estimate, which averages over the uncertainty in the attenuation for each beam aim. The data for each test or experiment are provided in Table V in the Appendix. As it was

TABLE III
ESTIMATED FAILURE FIT AND SDC FIT FOR CELL BEAMPATHS
AND OPTERON BEAMPATHS

	Failure FIT	SDC FIT
Cell Beampaths	172	7.2
95% CI	(92, 524)	(2.1, 26)
Opteron Beampaths	940	119
95% CI	(306, 2934)	(30, 453)

not possible to test one Cell in a Triblade without exposing the 390 second Cell in its beampath, the fluences include the exposure 391 gained when a Cell was running the idle test code while the 392 other Cell in its beampath was under test. 393

IV. RESULTS

A. Longevity of Hardware in the Beam, Post-Beam Testing and Root Cause Analysis of Permanently Failed Hardware

Some hardware experienced permanent failures relatively quickly upon exposure to the beam, while other hardware had greater longevity in the beam (see Table III). For example, QS22b on Triblade 2 was unable to boot after exposure to a corrected neutron fluence of *at most* $2.13 \times 10^9 \text{ neutrons/cm}^2$, while the lower Cell on QS22a on Triblade 4 remained operational after exposure to a corrected neutron fluence of *at least* $1.43 \times 10^{10} \text{ neutrons/cm}^2$. That said, Triblade 4 was tested with the one-inch beam diameter so it had less hardware in the beam than Triblade 2, which was tested with the two-inch beam diameter.

Following the beam testing, Triblades 1, 3 and 4 were tested in a production platform at LANL. (Triblade 2 had suffered damage through handling, and post-irradiation testing at LANL was not possible.) This testing used all of the applications from the beam testing, with the exception of the bottom Opteron in Triblade 4, which was not tested with the Opteron-only correlator code. Triblades 1, 3, and 4 each had a QS22 that would not boot. In addition, the QS22 in Triblade 1 that would boot could not communicate with the relevant Opteron.

After this post-irradiation testing was completed, all four Triblades were returned to IBM for root cause failure analysis. It was found that each of the 4 Triblades had a QS22 that had permanently failed. Further, these permanent failures were the result of voltage shorts in voltage regulators. Voltage regulators have been experimentally shown to experience single-event burnout (SEB) and single-event gate rupture (SEGR), both of which are destructive effects that can cause the system to be fully or partially unbiased, when exposed to thermal and fast neutrons [34]. The failed voltage regulators should not have been within the FWHM boundary of the beam unless the beam was mistargeted, so the neutron exposure they received should be less than that reported for the corresponding beam aims in Table II.

B. Failure Data

Each experiment was categorized as having one of two end states: 1) survival, meaning that the experiment ended when

the experimenter believed the application was still running or 2) failure, indicating that the application was no longer running at the end of the experiment, e.g., because of an application or a system crash. Since the output from the test applications appeared roughly every minute, it is possible that in some cases in which the system is deemed to have survived the experiment it had actually failed, but that failure was not detected before the experiment ended. Post-irradiation analysis showed that 79 of the 94 tests conducted on the Cells ended in failure, while all 14 tests conducted on the Opteron ended in failure.

Observed failures include application hangs, blades that spontaneously rebooted, and blades that became non-responsive. Investigation of the log data did not yield definitive root causes. Our hypothesis is that in most cases the hardware failed so completely and so quickly that no useful diagnostic information could be obtained.

C. Silent Data Corruption

In order to check for SDC, the computational test codes included a step in which the calculated answer was compared to the correct answer. Four SDCs were observed. Two SDCs occurred when a Cell was running a computational test code (one with VPIC and one with correlator) and two SDCs occurred when an Opteron was running the Opteron-only correlator test code. For the Cell beampaths, the median posterior probability that an error is an SDC rather than a failure (e.g., application or system crash) is 0.038 with 95% CI (0.011, 0.088), while for the Opteron beampaths it is 0.114 with 95% CI (0.035, 0.250). These estimates along with their corresponding uncertainty statements were obtained using standard Bayesian statistical methods for analyzing Binomial data [14]. The Opteron CI is wider because fewer experiments were conducted for the Opteron beampaths.

D. Failure FITs and SDC FITs for Cell and Opteron Beampaths

A statistical model that incorporated the upper and lower bounds on fluence until error (failure or SDC) and that accounted for the application used for each test, the Triblade under test, the beam aim (Cell beampaths or Opteron beampaths), and the beam diameter was fit to the experimental data [14]. All results presented below derive from this model and pertain to the conditions under which the experiments were conducted, with results likely to be obtained under other conditions less clear. Further, the results presented here are based on the experimental data collected at LANSCE and not on failures or SDCs observed in the Roadrunner platform. All results have been estimated via Markov Chain Monte Carlo [35].

Based on this modeling, Table II presents estimated failure FITs and SDC FITs for the Cell beampaths and the Opteron beampaths, along with 95% CIs that capture the uncertainty in the FIT estimates. These and all other FIT estimates presented in this work are based on one Cell idling while the other runs the VPIC test code (since it is most representative of a computational application that might be used in the field of those considered in our study) and the two-inch beam diameter.

TABLE IV
ESTIMATED FAILURE FIT AND SDC FIT FOR A TRIBLADE, A ROADRUNNER CU, AND ROADRUNNER

	Failure FIT	SDC FIT
Triblade	2.22×10^3	2.58×10^2
95% CI	(8.06×10^2 , 7.06×10^3)	(9.09×10^1 , 9.89×10^2)
Roadrunner CU	4.51×10^5	5.23×10^4
95% CI	(2.28×10^5 , 1.19×10^6)	(1.70×10^4 , 1.61×10^5)
Roadrunner	7.69×10^6	8.81×10^5
95% CI	(3.86×10^6 , 1.90×10^7)	(2.85×10^5 , 2.78×10^6)

They reflect the flux of neutrons in Los Alamos, NM that have energies greater than 10 MeV, which is estimated to be normal with mean 0.019 neutrons/cm²/sec [36] and standard deviation of 4.4×10^{-4} neutrons/cm²/sec [3].

The FIT estimates and corresponding uncertainty intervals are calculated by standard Bayesian analysis techniques. Specifically, a Monte Carlo procedure is used that repeatedly generates values of parameters from their posterior distribution based on the statistical model described, each time calculating FITs based on the generated parameter values. That is, the expected number of failures in 10^9 hours will be different for different parameter values. Thus, this procedure reflects the uncertainty in FIT due to the uncertainty in the unknown model parameters and the uncertainty in the amount of neutron exposure to the hardware under test. The FIT estimate is taken to be the median of the FITs calculated from the generated parameter values, while the 0.025 and 0.975 quantiles are used for the bounds of the 95% CIs. See [14] for details.

From the results in Table IV, the Cell beampaths are less susceptible to neutron-induced errors than the Opteron beampaths. Care must be taken in interpreting this result since these beampaths include hardware in addition to the microprocessor that was running applications during the testing. That is, the values in the Table II cannot be interpreted as reflecting only the Cell and Opteron microprocessors. In particular, a small amount of the Opteron memory was in the beam when the Cells were being tested, with more exposure resulting when using the two-inch beam diameter as opposed to the one-inch beam diameter. Similarly, a small amount of Cell memory was in the beam when testing one of the Opteron, but the Opteron in a particular Triblade were tested after the Cells in that Triblade were tested. Using the two-inch beam diameter versus the one-inch beam diameter does not significantly change the hazard rate or instantaneous error rate (see Section IV.E), suggesting that any resulting effects in the Opteron memory are not likely to be substantial.

Thus, while this study underscores that there is almost certainly a difference in neutron susceptibility between the hardware in the Opteron beampaths and the hardware in the Cell beampaths, identifying the source of this difference with certainty is not possible. Since all of the hardware in the beampaths of each of the processors was irradiated, it could reflect neutron interactions with this hardware rather than the processors themselves. Assuming that most if not all neutron effects occurred in the processors it could reflect their process technologies (the Cell is 65nm SOI and the Opteron is 90nm

534 SOI), transistor counts, caches sizes, numbers of susceptible
535 states, architectural vulnerability factors [37], [38], architec-
536 tures (the Cell architecture is somewhat simpler than that of the
537 Opteron) or some other cause.

538 *E. Effects of Application, Beam Aim, Beam Diameter, and* 539 *Triblade Under Test on the Error Rate*

540 Based on the results of the model described in Section IV-D,
541 the paragraphs below discuss the effects of increasing exposure
542 to the beam, beam aim, Triblade under test, application used
543 for the test, and beam diameter on the hazard rate, i.e., the
544 instantaneous error (failure and SDC) rate of the hardware
545 under test.

546 The baseline hazard rate appears to be close to constant,
547 suggesting that the instantaneous error rate likely does not vary
548 much with increasing exposure to the beam for the exposures
549 observed in this study. Therefore, it is likely that sensitivity
550 to neutrons does not change with increasing dose accumula-
551 tion and in-field usage should have roughly constant neutron-
552 induced error rates.

553 The posterior probability that the beam aim (Cell beampaths
554 or Opteron beampaths) affects the hazard rate is 1.0, meaning
555 that there is most certainly a difference in neutron sensitivity
556 between the hardware in the Cell beampaths and the hardware
557 in the Opteron beampaths. With the Opteron beampaths, the
558 median multiplier to the hazard rate is 5.884 with 95% CI
559 (2.749, 11.753), meaning that errors are roughly six times
560 more frequent with the Opteron beampaths than with the Cell
561 beampaths.

562 There is a relationship between the Triblade under test
563 and the beam diameter used for the testing. Triblade 3 was
564 tested using the two-inch beam diameter and Triblade 4 was
565 tested using the one-inch beam diameter, while Triblade 1
566 was tested using both beam diameters. With a situation like
567 this, it can be difficult to determine which of Triblade under
568 test or beam diameter is more influential on the hazard rate.
569 That said, the posterior probability that one or both of Triblade
570 under test and beam diameter affects the hazard rate is 0.931,
571 and the results below suggest that Triblade under test is more
572 likely than beam diameter to affect the hazard rate.

573 The modeling results indicate a 0.897 posterior probability
574 that different Triblades under test experienced different sensi-
575 tivities to the beam. The posterior median relative difference
576 in hazard rate for two randomly-selected Triblades is 1.357
577 with 95% CI (1.000, 5.049). Thus, this test data suggests
578 that process-variation-based differences in neutron sensitivity
579 may exist. However, more Triblades would need to be tested
580 and/or more time would need to be spent under test to fully
581 investigate the implications of process-variation-based neutron
582 sensitivities.

583 Beam diameter (one-inch versus two-inch) has a 0.198 pos-
584 terior probability of affecting the hazard rate, suggesting that
585 beam diameter did not have much if any impact on the hazard
586 rate. This implies that most of the sensitive hardware likely lies
587 within the one-inch beam diameter.

588 For the most part, the application being run did not affect
589 the hazard rate. The largest effect on the hazard rate is for

hybrid Linpack, with a 0.417 posterior probability of having
a hazard rate different from that of the idle condition. Its
median multiplicative effect on the hazard rate is 1.000, with
95% CI (1.000, 2.545). Therefore, the error sensitivity did not
have much application dependence. This result is consistent
with related findings in [5] and the confidence limits presented
in [4].

There are a number of possible explanations for this result.
First, the operating system, which executed in all tests whether
an application was executing or not, might be overshadowing
the effect of the application on the hardware sensitivity to
neutrons. In [39] results from [16] are used to indicate that the
proton cross-section for the Pentium II and MMX microproces-
sors was two to three orders of magnitude larger when tested
with Windows operating system than without. Since definitive
root causes for observed failures could not be determined, it
could be that enough failures resulted from OS tasks rather
than application tasks that it is not possible to distinguish large
differences among the applications. Second, the applications
chosen here may have similar neutron sensitivities, which other
applications might not share. Further study with more appli-
cations with different programming and computing patterns
would be useful. To better understand the extent to which
failures derive from OS tasks, the testing could be performed
with the applications running on the processors under test, but
without an OS.

566 *F. Projected Failure and SDC Rates for Roadrunner*

Roadrunner is composed of 17 connected units (CU), each
of which includes 180 Triblades that are used for computation.
The experimental results can further be used to estimate failure
FITs and SDC FITs and corresponding 95% CIs for a single
Triblade, for the 180 Triblades in a CU, and for all of the
Triblades in the Roadrunner platform (17 CUs); Table IV
provides these values.

These results do not reflect the neutron sensitivity of all of
the hardware in a Triblade, as they only include the hardware
in the Cell and Opteron beampaths. For the Triblade values
they assume that errors in the hardware in the different beam-
paths occur independently, while the CU values further assume
independence of errors in the Triblades within a CU and the
Roadrunner values assume independence of all Triblades within
Roadrunner. See Section IV-D for additional assumptions un-
derlying these FIT estimates.

Table IV indicates that for a Triblade, Roadrunner CU, and
Roadrunner the failure FIT estimate is roughly an order of
magnitude larger than the SDC FIT estimate. Roadrunner is
estimated to experience one cosmic-ray-neutron-induced fail-
ure roughly every 130 hours of operation and one cosmic-ray-
neutron-induced SDC roughly every 1100 hours of operation.

The effect of any SDCs on calculations performed on Road-
runner is likely to be small since the results of many cal-
culations are typically combined to produce a final result,
thus mitigating the effect of an SDC in any one of the un-
derlying calculations. Specifically, verification and validation
efforts involve parameter studies that enable errors bars to be
investigated and better understood, with a suite of calculations,

TABLE V
EXPERIMENTAL DATA

Record	Hardware Tested	Application	SDC	Fluence A	Fluence B
1	cell: 3b-low	varied	0	4.81×10^8	1.51×10^9
2	cell: 3b-low	varied	0	1.21×10^8	5.04×10^8
3	cell: 3b-low	varied	0	8.21×10^7	3.10×10^8
4	cell: 3b-low	corr	0	3.56×10^6	1.12×10^8
5	cell: 3b-low	corr	0	7.98×10^6	1.22×10^8
6	cell: 3b-low	varied	0	4.23×10^7	1.90×10^8
7	cell: 3b-upp	varied	0	1.37×10^8	3.59×10^8
8	cell: 3b-upp	varied	0	1.42×10^8	3.73×10^8
9	cell: 3b-upp	varied	0	6.71×10^7	3.94×10^8
10	cell: 3b-upp	varied	0	5.09×10^8	1.16×10^9
11	cell: 3b-upp	vpic	0	6.78×10^7	2.13×10^8
12	cell: 3b-upp	hpl	0	0.00	9.34×10^7
13	cell: 3b-upp	hpl	0	1.31×10^8	3.52×10^8
14	cell: 3b-upp	hpl	0	6.72×10^7	2.21×10^8
15	cell: 3b-upp	hpl	0	4.87×10^6	3.07×10^8
16	cell: 3b-upp	corr	0	4.70×10^7	1.77×10^8
17	cell: 3b-upp	corr	0	4.30×10^8	9.44×10^8
18	cell: 3b-upp	corr	0	1.19×10^8	3.30×10^8
19	cell: 3b-upp	corr	0	8.56×10^8	Inf
20	cell: 3b-upp	vpic	0	2.99×10^7	1.21×10^8
21	cell: 3b-upp	vpic	0	4.52×10^7	9.05×10^7
22	cell: 3b-upp	vpic	0	5.31×10^7	1.06×10^8
23	cell: 3b-upp	vpic	0	2.00×10^8	4.38×10^8
24	cell: 3b-upp	vpic	0	9.51×10^8	2.05×10^9
25	cell: 3b-upp	idle	0	3.49×10^8	7.48×10^8
26	cell: 3b-upp	int_add	0	2.19×10^8	4.99×10^8
27	cell: 3b-upp	cg	0	7.26×10^7	Inf
28	cell: 3b-upp	cg	0	9.66×10^7	3.86×10^8
29	cell: 3b-upp	corr	0	7.06×10^7	3.11×10^8
30	cell: 3b-upp	idle	0	6.69×10^8	Inf
31	cell: 3b-upp	int_add	0	6.49×10^8	Inf
32	cell: 3b-upp	vpic	0	1.27×10^8	3.74×10^8
33	cell: 3b-upp	vpic	1	2.14×10^7	6.65×10^7
34	cell: 3b-upp	vpic	0	2.20×10^8	Inf
35	cell: 3a-upp	hpl	0	3.04×10^8	8.02×10^8
36	cell: 3a-upp	hpl	0	3.56×10^7	2.73×10^8
37	cell: 3a-upp	hpl	0	0.00	1.32×10^8
38	cell: 3a-upp	hpl	0	1.87×10^7	1.54×10^8
39	cell: 3a-upp	hpl	0	1.58×10^8	Inf
40	cell: 3a-upp	int_add	0	4.50×10^8	1.11×10^9
41	cell: 3a-upp	int_add	0	1.98×10^8	4.54×10^8
42	cell: 3a-upp	int_add	0	9.42×10^7	Inf
43	cell: 3a-upp	hpl	0	2.64×10^7	1.51×10^8
44	cell: 3a-upp	hpl	0	1.35×10^8	3.40×10^8
45	cell: 3a-upp	hpl	0	3.23×10^7	8.94×10^7
46	cell: 3a-upp	hpl	0	4.06×10^8	9.22×10^8
47	cell: 1a-upp	corr	0	9.18×10^6	8.51×10^7

TABLE V
(Continued). EXPERIMENTAL DATA

Record	Hardware Tested	Application	SDC	Fluence A	Fluence B
48	cell: 1a-upp	corr	0	2.91×10^8	6.63×10^8
49	cell: 1a-upp	corr	1	1.31×10^8	2.95×10^8
50	cell: 1a-upp	corr	0	1.55×10^8	7.63×10^8
51	cell: 1a-upp	cg	0	2.67×10^8	5.95×10^8
52	cell: 1a-upp	cg	0	4.84×10^7	1.92×10^8
53	opt: 1-top	corr	0	4.35×10^7	1.26×10^8
54	opt: 1-top	corr	1	4.92×10^6	4.26×10^7
55	opt: 1-top	corr	0	0.00	5.01×10^7
56	opt: 1-top	corr	0	1.53×10^6	4.95×10^7
57	cell: 4b-upp	corr	0	4.68×10^8	9.80×10^8
58	cell: 4b-upp	corr	0	5.83×10^8	Inf
59	cell: 4a-upp	cg	0	4.56×10^8	Inf
60	cell: 4a-upp	corr	0	2.69×10^8	6.16×10^8
61	cell: 4a-upp	corr	0	2.16×10^8	Inf
62	cell: 4a-upp	int_add	0	2.60×10^8	5.60×10^8
63	cell: 4a-upp	hpl	0	8.29×10^8	Inf
64	cell: 4a-upp	idle	0	2.56×10^7	2.00×10^8
65	cell: 4a-upp	idle	0	2.00×10^8	4.68×10^8
66	cell: 4a-upp	idle	0	4.14×10^7	1.73×10^8
67	cell: 4a-upp	cg	0	1.31×10^8	3.49×10^8
68	cell: 4a-upp	cg	0	3.34×10^7	1.23×10^8
69	cell: 4b-low	vpic	0	7.15×10^8	Inf
70	cell: 4b-low	cg	0	6.52×10^8	1.83×10^9
71	cell: 4b-low	corr	0	2.03×10^8	1.29×10^9
72	cell: 4b-low	corr	0	1.74×10^8	6.54×10^8
73	cell: 4b-low	corr	0	1.75×10^8	5.21×10^8
74	cell: 4b-low	int_add	0	5.36×10^8	1.47×10^9
75	cell: 4b-low	int_add	0	7.15×10^8	1.93×10^9
76	cell: 4b-low	hpl	0	4.99×10^7	1.75×10^8
77	cell: 4b-low	hpl	0	5.24×10^8	1.43×10^9
78	cell: 4b-low	idle	0	3.56×10^8	9.83×10^8
79	cell: 4b-low	vpic	0	7.94×10^8	2.17×10^9
80	cell: 4b-low	corr	0	4.24×10^7	4.67×10^8
81	cell: 4b-low	corr	0	1.82×10^8	6.01×10^8
82	cell: 4b-low	corr	0	2.20×10^8	6.62×10^8
83	cell: 4b-low	cg	0	6.51×10^8	2.07×10^9
84	cell: 4b-low	int_add	0	5.01×10^8	1.37×10^9
85	cell: 4b-low	hpl	0	4.33×10^7	1.99×10^8
86	cell: 4b-low	hpl	0	8.90×10^6	5.87×10^7
87	cell: 4b-low	hpl	0	5.97×10^7	2.33×10^8
88	cell: 4b-low	hpl	0	4.05×10^8	Inf
89	cell: 4b-low	vpic	0	2.61×10^8	Inf
90	cell: 4b-low	vpic	0	1.66×10^8	6.69×10^8
91	cell: 4b-low	vpic	0	3.56×10^6	1.08×10^8
92	cell: 4b-low	vpic	0	4.38×10^8	1.25×10^9
93	cell: 4b-low	int_add	0	8.48×10^8	Inf
94	cell: 4b-low	corr	0	1.87×10^9	Inf

TABLE V
(Continued). EXPERIMENTAL DATA

Record	Hardware Tested	Application	SDC	Fluence A	Fluence B
95	cell: 4b-low	cg	0	7.31×10^8	1.96×10^9
96	cell: 4b-low	cg	0	3.52×10^7	1.13×10^8
97	cell: 4b-low	cg	0	1.03×10^8	3.00×10^8
98	cell: 4b-low	cg	0	1.24×10^8	4.05×10^8
99	cell: 4b-low	cg	0	1.71×10^8	4.89×10^8
100	cell: 4b-low	cg	0	1.33×10^8	4.02×10^8
101	cell: 4b-low	cg	0	8.70×10^8	2.75×10^9
102	opt: 4-low	corr	0	1.72×10^7	1.58×10^8
103	opt: 4-low	idle	0	1.41×10^8	3.70×10^8
104	opt: 4-low	idle	0	6.00×10^7	1.57×10^8
105	opt: 4-low	corr	0	0.00	1.98×10^8
106	opt: 4-low	corr	0	3.80×10^6	1.59×10^8
107	opt: 4-low	idle	0	3.84×10^7	1.01×10^8
108	opt: 1-low	corr	0	7.04×10^6	6.12×10^7
109	opt: 1-low	idle	0	3.42×10^6	4.05×10^7
110	opt: 1-low	idle	0	3.24×10^7	1.75×10^8
111	opt: 1-low	corr	1	4.94×10^7	1.62×10^8
112	opt: 1-low	corr	0	3.67×10^7	1.75×10^8
113	opt: 1-low	corr	0	1.88×10^7	3.13×10^8

including some used to investigate error bars, used for decision making.

V. CONCLUSION

Replicates of two microprocessors, the IBM PowerXCell 8i and the AMD Opteron 2210 HE, along with the hardware in their respective beampaths, were tested at LANSCE for neutron sensitivities. These tests indicated that both microprocessor beampaths were susceptible to neutron-induced errors and that the Opteron beampaths were more sensitive to neutrons than the Cell beampaths as evidenced by the failure FIT and SDC FIT estimated for each of these beampaths. The data further provided some evidence for process-variation-based neutron sensitivity differences. Little application-based neutron sensitivity differences were found, with hybrid Linpack most likely to lead to a somewhat elevated hazard rate. The results suggest that failures, e.g., application and system crashes, occur roughly an order of magnitude more often than SDCs for the Triblades under test and for the Roadrunner platform that leverages them for computation.

APPENDIX EXPERIMENTAL DATA

Table V provides the experimental conditions pertaining to and data collected for each of the 113 errors analyzed for the results presented here. These errors include 109 experiments that ended with a failure or an operator decision to terminate the experiment and 4 SDCs. The data includes the following columns: Record (which corresponds to the sequential order in

which errors were observed and of tests that an operator ended while the system remained operational); Hardware Tested (the Triblade and location on that Triblade at which the beam was aimed; in the case of Cells running a computational code, it also provides which Cell was running the computational code, i.e., cell: 3a-upp, means Triblade 3 was in the beam, with the beam aimed at the upper Cells with the upper Cell in QS22a running a computational application); Application (the test code that was run prior to the error (crash or SDC) denoted as follows: hpl (hybrid Linpack), corr (correlator), cg (conjugate gradient), vplic (VPIC), integer; adder (int_add), varied (varied), and idle (idle)); SDC (a value of 1 indicates that an SDC occurred, with a 0 if otherwise), Fluence A (posterior mean of the lower bound for the neutron fluence for neutrons with energies above 10 MeV accumulated at the processor under test until error), and Fluence B (posterior mean of the upper bound for the neutron fluence for neutrons with energies above 10 MeV accumulated at the processor under test until error, with a value of “Inf” indicating that the operator decided to terminate the experiment prior to an error occurring).

ACKNOWLEDGMENT

The authors would like to thank J. Abeyta, C. Alexander, B. Bergen, A. Borrett, H. Brandt, J. Campa, R. Cardon, N. De Bardeleben, T. Fairbanks, P. Fields, A. Gibson, G. Grider, J. Loncaric, P. Lujan, A. Malin, F. Marshall, A. Montoya, J. Morrison, A. Shewmaker, M. Vigil, B. Villa, S. Wender, A. White, and C. Wright. The authors apologize for any inadvertent omissions from this list. The authors further thank the reviewers for their insightful comments.

The U.S. Government has rights to use, reproduce, and distribute this information. The public may copy and use this information without charge, provided that this Notice and any statement of authorship are reproduced on all copies. Neither the Government nor LANS makes any warranty, express or implied, or assumes any liability or responsibility for the use of this information. The Los Alamos National Laboratory strongly supports academic freedom and a researcher’s right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness. This paper is published under LA-UR-12-20226. Copyright copyright 2010 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained by sending a request to pubs-permissions@ieee.org.

REFERENCES

- [1] R. Baumann, “Radiation induced soft errors in advanced semiconductor technologies,” *IEEE Trans. Device Mater. Rel.*, vol. 5, no. 3, pp. 305–316, Sep. 2005.
- [2] C. Constantinescu, “Neutron SER characterization of microprocessors,” in *Proc. Int. Conf. Dependable Syst. Netw.*, 2005, pp. 754–759.
- [3] S. Michalak, K. Harris, N. Hengartner, B. Takala, and S. Wender, “Predicting the number of fatal soft errors in Los Alamos National Laboratory’s ASC Q supercomputer,” *IEEE Trans. Device Mater. Rel.*, vol. 5, no. 3, pp. 329–335, Sep. 2005.
- [4] H. Ando, R. Kan, Y. Tosaka, K. Takahisa, and K. Hatanaka, “Validation of hardware error recovery mechanisms for the SPARC64 V microprocessor,” in *Proc. IEEE Int. Conf. Dependable Syst. Netw.*, 2008, pp. 62–69.

- [5] S. Rao, T. Hong, P. Sanda, J. Ackaret, A. Barrera, J. Yanez, S. Mitra, J. Kellington, and R. McBeth, "Examining workload dependence of soft error rates," in *Proc. SELSE*, 2008.
 - [6] P. Sanda, J. Kellington, P. Kudva, R. Kalla, R. McBeth, J. Ackaret, R. Lockwood, J. Schumann, and C. Jones, "Soft_error resilience of the IBM POWER6 processor," *IBM J. Res. Develop.*, vol. 52, no. 3, pp. 275–284, 2008.
 - [7] T. Hong, S. Michalak, T. Graves, J. Ackaret, and S. Rao, "Neutron beam irradiation study of workload dependence of SER in a microprocessor," in *Proc. SELSE*, 2009, pp. 1–5.
 - [8] F. Wrobel, J. Gasiot, and F. Saigne, "Hafnium and uranium contributions to soft error rate at ground level," *IEEE Trans. Nucl. Sci.*, vol. 55, no. 6, pp. 3141–3145, Dec. 2008.
 - [9] J. Autran, D. Munteanu, P. Roche, G. Gasiot, S. Martinie, S. Uznanski, S. Sauze, S. Semikh, E. Yakushev, S. Rozov, P. Loaiza, G. Warot, and M. Zampaolo, "Soft_errors induced by terrestrial neutrons and natural alpha_particle emitters in advanced memory circuits at ground level," *Microelectron. Reliab.*, vol. 50, no. 9_11, pp. 1822–1831, 2010.
 - [10] Z. Kalbarczyk, R. Iyer, and L. Wang, "Application fault tolerance with armor middleware," *IEEE Internet Comput.*, vol. 9, no. 2, pp. 28–37, Mar./Apr. 2005.
 - [11] D. Katz, *Application-Based Fault Tolerance for Spaceborne Applications*. [Online]. Available: <http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/10574/1/02-2631.pdf>
 - [12] K. Koch, *Roadrunner Platform Overview*, 2008. [Online]. Available: <http://www.lanl.gov/orgs/hpc/roadrunner/pdfs/Koch%20-%20Roadrunner%20Overview/RR%20Seminar%20-%20System%20Overview.pdf>
 - [13] H. Meuer, *31st TOP500 List Topped by First_Ever Petaflop/s Supercomputer Scientific Computing*, 2008. [Online]. Available: <http://www.scientificcomputing.com/31st-TOP500-List-Topped-by-First-ever-Petaflops-Supercomputer.aspx>
 - [14] C. Storlie, S. Michalak, H. Quinn, A. DuBois, and D. DuBois, "A Bayesian survival analysis of neutron-induced errors in high performance computing hardware," Los Alamos Nat. Lab., Los Alamos, NM, 2011.
 - [15] S. Michalak, A. DuBois, C. Storlie, H. Quinn, W. Rust, D. DuBois, D. Modl, A. Manuzzato, and S. Blanchard, "Neutron beam testing of high performance computing hardware," in *Proc. IEEE Nucl. Space Radiation Effects Conf. Radiation Effects Data Workshop*, 2011, pp. 1–8.
 - [16] D. Hiemstra, A. Baril, and M. Dettwiler, "Single event upset characterization of the Pentium (R) MMX and Pentium (R) II microprocessors using proton irradiation," *IEEE Trans. Nucl. Sci.*, vol. 46, no. 6, pp. 1453–1460, Dec. 1999.
 - [17] D. Hiemstra and A. Baril, "Single event upset characterization of the Pentium (R) MMX and Celeron (R) microprocessors using proton irradiation," in *Proc. IEEE Radiation Effects Data Workshop*, 2000, pp. 39–44.
 - [18] D. Hiemstra, S. Yu, and M. Pop, "Single event upset characterization of the Pentium (R) MMX and Low Power Pentium (R) MMX microprocessors using proton irradiation," in *Proc. IEEE Radiation Effects Data Workshop*, 2001, pp. 32–37.
 - [19] D. Hiemstra, S. Yu, and M. Pop, "Single event upset characterization of the Pentium 4, Pentium III and Low Power Pentium MMX microprocessors using proton irradiation," in *Proc. IEEE Radiation Effects Data Workshop*, 2002, pp. 51–57.
 - [20] G. Swift, F. Farmanesh, F. Guertin, F. Irom, and D. Millward, "Single_event upset in the PowerPC750 microprocessor," *IEEE Trans. Nucl. Sci.*, vol. 48, no. 6, pp. 1822–1827, Dec. 2001.
 - [21] F. Irom, F. Farmanesh, and C. Kouba, "Single_event upset and scaling trends in new generation of the commercial SOI PowerPC microprocessors," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 6, pp. 3563–3568, Dec. 2006.
 - [22] F. Irom, T. Miyahira, D. Nguyen, I. Jun, and E. Normand, "Results of recent 14 MeV neutron single event effects measurements conducted by the Jet Propulsion Laboratory," in *Proc. IEEE Radiation Effects Data Workshop*, 2007, pp. 141–145.
 - [23] A. Dixit, R. Heald, and A. Wood, "Trends from ten years of soft error experimentation," in *Proc. SELSE*, 2009, pp. 1–4.
 - [24] A. Dixit and A. Wood, "The impact of new technology on soft error rates," in *Proc. SELSE*, 2011, pp. 486–492.
 - [25] M. Gschwind, P. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, and T. Yamazaki, *A Novel SIMD Architecture for the Cell Heterogeneous Chip Multiprocessor*, 2005. [Online]. Available: http://www.hotchips.org/archives/hc17/2_Mon/HC17.S1/HC17.S1T1.pdf
 - [26] AMD, *The AMD Opteron Processor for High-End Embedded Designs*, 2008. [Online]. Available: http://www.amd.com/es/Documents/42350B_Operon_PB.pdf
 - [27] IBM, *Bladecenter H Chassis*. [Online]. Available: <http://www-03.ibm.com/systems/bladecenter/hardware/chassis/bladeh/>
 - [28] *Virtex-II Platform FPGAs : Complete Data Sheet (DS031)*. [Online]. Available: <http://www.xilinx.com>
 - [29] H. Quinn, P. Graham, M. Wirthlin, B. Pratt, K. Morgan, M. Caffrey, and J. Krone, "A test methodology for determining space readiness of Xilinx SRAM_based FPGA devices and designs," *IEEE Trans. Instrum. Meas.*, vol. 58, no. 10, pp. 3380–3395, Oct. 2009.
 - [30] M. Kistler, J. Gunnels, D. Brokenshire, and B. Benton, "Programming the Linpack benchmark for Roadrunner," *IBM J. Res. Develop.*, vol. 53, no. 5, pp. 736–746, Sep. 2009.
 - [31] A. Petitet, R. C. Whaley, J. Dongarra, and A. Cleary, *HPL_A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers*. [Online]. Available: <http://netlib.org/benchmark/hpl/>
 - [32] A. DuBois, C. Conner, S. Michalak, G. Taylor, and D. DuBois, "Application of the IBM cell processor to real-time cross-correlation of a large antenna array radio telescope," Los Alamos Nat. Lab., Los Alamos, NM, Tech. Rep. LA-UR-09-03483, 2009.
 - [33] K. J. Bowers, B. J. Albright, L. Yin, W. Daughton, V. Roytershteyn, B. Bergen, and T. J. T. Kwan, "Advances in petascale kinetic plasma simulation with VPIC and Roadrunner," *J. Phys., Conf. Ser.*, vol. 180, no. 1, p. 012055, 2009.
 - [34] A. Hands, P. Morris, K. Ryden, C. Dyer, P. Truscott, A. Chugg, and S. Parker, "Single event effects in power MOSFETS due to atmospheric and thermal neutrons," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 6, pp. 2687–2694, Dec. 2011.
 - [35] A. Gelman, J. Carlin, H. Stern, and D. Rubin, *Bayesian Data Analysis*, 2nd ed. London, U.K.: Chapman & Hall, 2004.
 - [36] *Flux Calculation*, 2008. [Online]. Available: <http://seutest.com/cgi-bin/FluxCalculator.cgi/>
 - [37] S. Mukherjee, C. Weaver, J. Emer, S. Reinhardt, and T. Austin, "A systematic methodology to compute the architectural vulnerability factors for a high-performance microprocessor," in *Proc. 36th Annu. Int. Symp. MICRO*, 2003, pp. 29–40.
 - [38] A. Biswas, P. Racunas, R. Cheveresan, J. Emer, S. Mukherjee, and R. Rangan, "Computing architectural vulnerability factors for address_based structures," in *Proc. 32nd ISCA*, 2005, pp. 532–543.
 - [39] F. Irom, "Guideline for ground radiation testing of microprocessors in the space radiation environment," Jet Propulsion Lab., Los Angeles, CA, JPL Pub. 08-13, 2008.
- Sarah E. Michalak**, photograph and biography not available at the time of publication.
- Andrew J. DuBois**, photograph and biography not available at the time of publication.
- Curtis B. Storlie**, photograph and biography not available at the time of publication.
- Heather M. Quinn**, photograph and biography not available at the time of publication.
- William N. Rust**, photograph and biography not available at the time of publication.
- David H. DuBois**, photograph and biography not available at the time of publication.
- David G. Modl**, photograph and biography not available at the time of publication.
- Andrea Manuzzato**, photograph and biography not available at the time of publication.
- Sean P. Blanchard**, photograph and biography not available at the time of publication.

AUTHOR QUERY

NO QUERY.

Assessment of the Impact of Cosmic-Ray-Induced Neutrons on Hardware in the Roadrunner Supercomputer

Sarah E. Michalak, Andrew J. DuBois, Curtis B. Storlie, Heather M. Quinn, William N. Rust, David H. DuBois, David G. Modl, Andrea Manuzzato, and Sean P. Blanchard

Abstract—Microprocessor-based systems are a common design for high-performance computing (HPC) platforms. In these systems, several thousands of microprocessors can participate in a single calculation that may take weeks or months to complete. When used in this manner, a fault in any of the microprocessors could cause the computation to crash or cause silent data corruption (SDC), i.e., computationally incorrect results that originate from an undetected fault. In recent years, neutron-induced effects in HPC hardware have been observed, and researchers have started to study how neutrons impact microprocessor-based computations. This paper presents results from an accelerated neutron-beam test focusing on two microprocessors used in Roadrunner, which is the first Petaflop supercomputer. Research questions of interest include whether the application running affects neutron susceptibility and whether different replicates of the hardware under test have different susceptibilities to neutrons. Estimated failures in time for crashes and for SDC are presented for the hardware under test, for the Triblade servers used for computation in Roadrunner, and for Roadrunner.

Index Terms—Failures in time (FIT), neutron-beam testing, silent data corruption (SDC), single-event effect, soft error.

I. INTRODUCTION

LARGE-SCALE scientific computations are frequently performed on high-performance computing (HPC) platforms. These computations can use thousands of processors and run for weeks to months. Many HPC platforms use commercial off-the-shelf (COTS) microprocessors, as opposed to radiation-hardened devices, for such computation. Neutron-induced effects in COTS microprocessors include single-event upsets (SEUs) in the caches, register files, pipeline registers, and memory; single-event transients (SETs) in functional units;

and single-event functional interrupts (SEFIs) in control logic. The results of such neutron-induced effects can include failures (e.g., system and application crashes) and silent data corruption (SDC), which occurs when an undetected error causes the system to deliver computationally incorrect results. Neutrons have been implicated in crashes and SDC in different architectures [1]–[7]. While alpha particles can lead to similar issues [8], [9], this paper focuses on the effects of neutrons since it is concerned with the experience of systems located at Los Alamos National Laboratory, which is at high elevation (7200 ft) and hence experiences a higher neutron flux than that at sea level.

Because SEUs, which include single-bit upsets (SBUs) and multi-bit upsets (MBUs), are increasingly noticeable in terrestrial applications, COTS microprocessor designers and system designers often include some protection from SEUs. These protections include error-correcting codes (ECC), bit interleaving in caches, and parity checks. For computations run on HPC platforms, software-level protections, such as checkpoint/restart, are also implemented. In checkpoint/restart, the calculation's state is periodically saved to hard disk so that a calculation can be restarted from the previous state if necessary. A second method that may be implemented is algorithm-based fault tolerance (ABFT), but it relies on specialized knowledge by the programmer or *ad hoc* optimization of the code by hand, e.g., [10], [11]. ABFT usually decreases performance, which is paramount in HPC systems, and is typically not used in HPC systems.

While these protections are useful, the failures and SDC that can result from neutron induced are not completely suppressed. HPC platforms used for scientific computation are particularly sensitive to neutron-induced faults due to their large size and the availability requirements of the applications that run on them. An HPC platform can contain thousands of replicates of a particular microprocessor or other sensitive device, making these large platforms more likely to experience the effects of cosmic-ray-induced neutrons than smaller systems.

With a calculation using many processors, it is often the case that if a single microprocessor crashes the entire calculation will be stopped, the previous checkpoint loaded, and the calculation restarted. For these situations, the application runtime is increased each time the calculation is restarted from a checkpoint since all of the runtime between the checkpoint and the crash is lost. Furthermore, SDC can be difficult to detect. Therefore, in HPC platforms neutron-induced effects are

Manuscript received September 23, 2011; revised January 26, 2012; accepted February 27, 2012. This work was supported in part by the Los Alamos National Security, LLC under Contract DE-AC52-06NA25396 and in part by the U.S. Department of Energy.

S. E. Michalak, C. B. Storlie, and W. N. Rust are with the Statistical Sciences Group, Los Alamos National Laboratory, Los Alamos, NM 87544 USA (e-mail: michalak@lanl.gov; storlie@lanl.gov; wnr@lanl.gov).

A. J. DuBois, D. H. DuBois, D. G. Modl, and S. P. Blanchard are with the Systems Engineering and Integration Group, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: ajd@lanl.gov; dhd@lanl.gov; digem@lanl.gov; seanb@lanl.gov).

H. M. Quinn is with the Space Data Systems Group, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: hquinn@lanl.gov).

A. Manuzzato was with Università degli Studi di Padova, 35122 Padova, Italy (e-mail: andrea.manuzzato@iee.org).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TDMR.2012.2192736

82 of concern since (1) system crashes affect application runtimes
83 and (2) SDC in scientific applications may lead to incorrect
84 scientific conclusions.

85 This paper presents results from neutron-beam testing of
86 hardware identical to that used in Roadrunner [12], the first
87 Petaflop system [13]. The hardware was tested while running
88 different applications including some used for scientific re-
89 search. The structure of this paper is as follows. Section II
90 discusses related microprocessor studies. Section III presents
91 the test setup, with Section IV detailing the results. Section V
92 offers conclusions from this work. The statistical methods for
93 the data are described in [14], while this work augments and
94 complements [15] by presenting data pertaining to permanent
95 failures and estimated failures in time (failures in 10^9 hours
96 of operation or FIT) for failures (e.g., application and system
97 crashes) and for SDC for the microprocessors tested and the
98 other hardware in their beampaths, for the Triblade server
99 used for computation in the Roadrunner platform, and for the
100 Roadrunner platform itself.

101 II. RELATED WORK

102 There is more than a decade's worth of static test data on
103 microprocessors [16]–[20], with a number of recent publica-
104 tions addressing more modern microprocessors with reduced
105 feature sizes or multiple processing cores [4], [21]–[24]. While
106 static testing is often the basis for error rate calculations, it
107 can be difficult to translate the errors from static tests into
108 dynamic error rates that reflect the field experience of real-
109 world systems. Determining the overall effect of radiation on
110 microprocessors is not simple, as faults in a system can remain
111 dormant for several thousands of clock cycles before triggering
112 an error or can be masked. In addition, the operating system and
113 any software being used can create noise in the system, making
114 it difficult to determine the cause of system crashes.

115 Several studies performed dynamic testing similar to that
116 performed here. These include [4] (SPARC64 V microproces-
117 sor), [5] (IBM POWER6 and Intel Core2 5160 Xeon Wood-
118 crest microprocessors), [7] (Intel Core2 5160 Xeon Woodcrest
119 microprocessor), [6] (IBM POWER6 microprocessor), and [2]
120 (Intel Itanium processor). All of these studies except [5] explic-
121 itly report observing SDC or events that could lead to SDC.
122 Further, [4], [5], and [7] studied whether different applications
123 led to differing susceptibilities to neutrons. [5] did not establish
124 differing susceptibilities, and while point estimates of logic der-
125 ating factors provided in [4] suggested some differences for the
126 applications used, 95% confidence intervals largely coincided.
127 [7] found that while the mean times to first indication of failure
128 (MTFIF) for some of the pairs of applications studied had
129 high probability of being different, the idle condition did not
130 have the highest MTFIF. Finally, [16]–[19] performed proton
131 testing of Pentium and Celeron microprocessors while running
132 consistency checks and a workload simulator, with both failures
133 and SDC observed.

134 III. TEST SETUP AND EXPERIMENTAL PROTOCOL

135 Hardware from Roadrunner was tested at Los Alamos Na-
136 tional Laboratory's (LANL) Los Alamos Neutron Science
137 Center (LANSCE) Irradiation of Chips and Electronics (ICE)

House in October 2009 to investigate the neutron susceptibility
138 of the two microprocessors used in Roadrunner along with the
139 hardware in their respective beampaths. Both microprocessors,
140 the IBM PowerXCell 8i (Cell) and the AMD Opteron 2210 HE,
141 have been commercially available. The test setup included test-
142 ing multiple replicates of the Cell microprocessor and Opteron
143 microprocessor while running different applications, including
144 some used for scientific computation.

145 The Cell microprocessors and Opteron microprocessors were
146 operated in the neutron beam in their field configuration in
147 a Triblade blade server. A Triblade [12] includes one IBM
148 LS21 blade, two IBM QS22 blades, and an expansion blade to
149 manage data traffic. The LS21 blade has two dual-core Opteron
150 2210 HE microprocessors, and the QS22 blades (QS22a and
151 QS22b) each have two Cell microprocessors.

152 The Cell is a 65nm silicon-on-insulator (SOI) microproces-
153 sor with 1 PowerPC processor element (PPE) that controls 8
154 synergistic processor elements (SPE); [25, p. 5] provides a
155 diagram of the Cell architecture. The 3.2 GHz PPE includes
156 a PowerPC processor unit (PPU) that is based on the PowerPC
157 architecture, a parity-protected 32 KB L1 data cache, a parity-
158 protected 32 KB L1 instruction cache, and a 512 KB L2
159 cache with ECC on data and parity on directory tags (which is
160 recoverable using redundant directories). Each 3.2 GHz SPE in-
161 cludes a synergistic processor unit (SPU) and an ECC-protected
162 256 KB dedicated non-caching local store. The QS22 blade
163 that housed the Cells during the testing included 8 GB of ECC
164 double data rate 2 (DDR2) dynamic random access memory
165 (DRAM).

166 The Opteron 2210 HE is a 1.8 GHz 90nm SOI dual-core
167 microprocessor. See [26, p. 2] for a diagram of the Opteron
168 2210 EE microprocessor, which has a design similar to the
169 Opteron 2210 HE tested at LANSCE. Each Opteron core has
170 an ECC-protected 64 KB L1 data cache, a parity-protected
171 64 KB L1 instruction cache, and an ECC-protected 1MB L2
172 cache. The LS21 blade that housed the Opteron 2210 HE for
173 the testing included 16 GB of ECC DDR2 DRAM.

174 A. Test Setup

175 There are two aspects of the test setup: the hardware test
176 setup and the software test setup. Both aspects were specifically
177 designed to mimic how the devices under test operate in the
178 field as part of the Roadrunner platform.

179 1) *Hardware Test Setup*: Due to their extreme size, most
180 HPC platforms need an efficient power, cooling and network
181 design, as the entire system might span thousands of square
182 feet. To this end, most platforms are designed to be housed
183 in server racks, with each rack housing multiple chassis. In
184 a blade-based platform such as Roadrunner, each chassis will
185 house multiple blades. In Roadrunner, the rack provides physi-
186 cal structure; the chassis provides a common interface to power,
187 network, and cooling; and the LS21 and QS22 compute blades
188 provide the compute infrastructure.

189 The hardware tested included four Triblades and a
190 BladeCenter-H (BC-H Type 8852) [27] chassis, which is de-
191 signed to house up to three Triblades. The testing required
192 additional hardware for system control and for neutron fluence
193

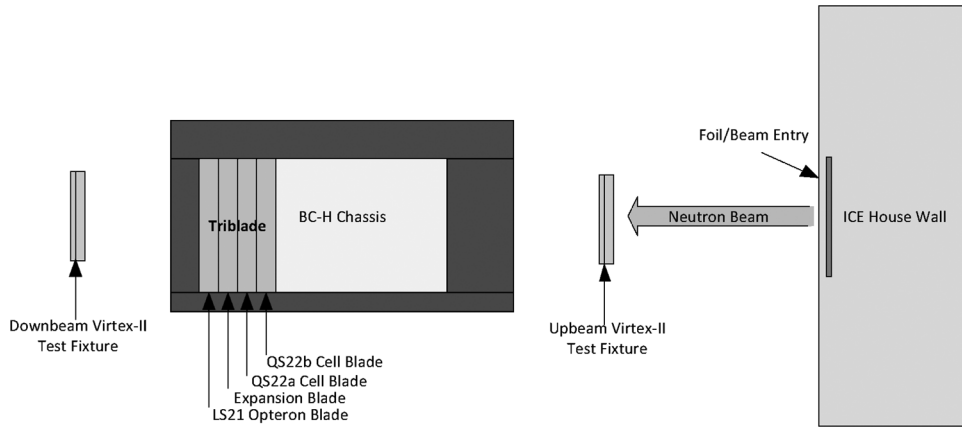


Fig. 1. Schematic of test setup with Triblade in the BC-H slot furthest from the beam source.

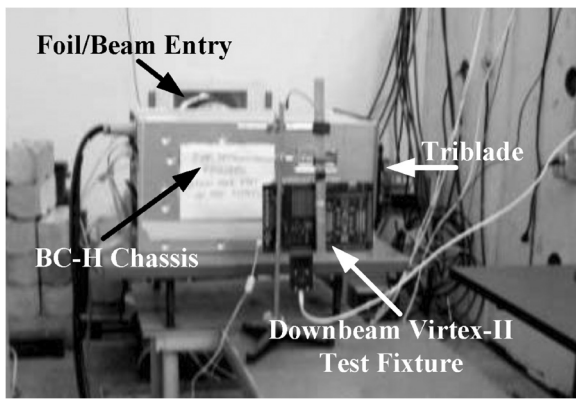


Fig. 2. Photo of test set-up.

194 exposure measurements. In many HPC platforms, a front-end
 195 node is used to manage the back-end compute nodes or blades,
 196 and likewise for this testing a front-end node was necessary to
 197 control the system under test. Specifically, it was used to boot
 198 the Triblades, start applications on the system under test and
 199 to monitor its health, all of which were performed manually
 200 by the experiment personnel. An IBM eServer X Series 336
 201 provided this capacity. This server was placed in the user area
 202 of the facility so that it would be protected from the neutron
 203 beam.

204 Since the hardware setup included more physical matter
 205 (chassis, metal enclosures, large heatsinks) than is typical, two
 206 Xilinx Virtex-II [28] test fixtures [29] were included in the test
 207 setup for calculating corrected neutron fluence exposures for
 208 the hardware under test. Specifically, one was placed upbeam
 209 of the BC-H and the other was placed downbeam of the BC-H.

210 Figs. 1 and 2 show the hardware test setup. Fig. 1 provides
 211 a diagram of the test facility and the hardware under test,
 212 including the point at which the beam enters the test facility,
 213 the Virtex-IIs, the BC-H chassis and the Triblade under test.
 214 Fig. 2 provides a complementary photograph with key aspects
 215 of the test setup labeled. The BC-H was oriented so that with a
 216 single Triblade under test, the beam first entered the QS22b,
 217 followed by the QS22a, the expansion blade, and the LS21
 218 respectively. For the testing, one Cell had a higher beam aim
 219 than the other, with “Upper Cell” denoting this Cell and “Lower
 220 Cell” denoting the Cell with the lower beam aim. For the two

Opertons in the LS21, “Upper Opteron ” and “Lower Opteron”
 have analogous meanings.

222
 223
 224
 225
 226
 227
 228
 229
 230
 231
 232
 233
 234
 235
 236
 237
 238
 239
 240
 241
 242
 243
 244
 245
 246
 247
 248
 249
 250
 251
 252
 253
 254
 255
 256
 257
 258
 259
 260
 261
 262

The Opteron test applications included an Opteron-only ver-
 sion of the correlator test code, idling, and running the Linux
 top command, which is considered an idle condition in the
 analyzes that follow. The Opteron-only correlator test code
 performs the multiply and accumulate described for the Cell
 above on a single Opteron core, with both cores on the Opteron
 under test running the code during the testing.

Each test code was designed so that it completed its work
 in roughly one minute. The software setup was instrumented

263 to run the test code continuously and return output data each
 264 time the test code completed. The output data included start and
 265 stop times, the application being run, the hardware running it,
 266 whether an SDC occurred and with the Cell idle code whether
 267 the Cells under test were still responding.

268 It should be noted that initial testing of the devices to
 269 determine the static sensitivity of the caches and registers to
 270 SEU was not undertaken. All of the FIT estimates and other
 271 results determined from this study are based on the described
 272 dynamic system use.

273 B. Experimental Procedure

274 For a given experiment, a single Cell or Opteron was config-
 275 ured to run the desired application while the beam was aimed
 276 so that it irradiated all of the hardware in the Triblade and the
 277 BC-H in that microprocessor's beampath. With two QS22s in
 278 a Triblade, when a Cell in one QS22 is running an application,
 279 the corresponding Cell in the other QS22 is in the beampath.
 280 This second Cell in the beampath was set to run the Cell idle
 281 test code. Since the beam irradiated a cylindrical volume within
 282 the Triblade under test and the BC-H, certain attribution of an
 283 error to the Cells or Opteron in the beampath is not possible.
 284 In particular, other hardware in the beampath or hardware that
 285 was affected by scatter could be the cause of an observed error.
 286 Errors could also be the result of causes external to the beam,
 287 but this is much less likely.

288 The experimental protocol was to start the required test
 289 application on the appropriate microprocessor while the beam
 290 was off. Once the test application was observed to be operating
 291 properly (e.g., the test code had produced one or more output
 292 lines), the beam was started. The experiment continued until
 293 a state of system inoperability (e.g., a system or application
 294 crash) was reached or until sufficient time had elapsed. The
 295 beam was then turned off, data pertaining to neutron fluence
 296 exposure of the system under test were collected, and the sys-
 297 tem was rebooted before beginning the next test. For the Cells,
 298 the test procedure was to cycle through the test applications on a
 299 particular Cell, typically until it became inoperable. Repeating
 300 each test code periodically permits investigation of any aging or
 301 dose-related effects resulting from increasing exposure to the
 302 beam. The procedure for the Optérons, which received much
 303 less testing, was to use the Opteron-only correlator code and
 304 possibly an idle condition (idling or running the Linux top
 305 command). Functionality of the Opteron while it was idling or
 306 running the Linux top command was assessed by ascertaining
 307 its continued responsiveness.

308 In all, 113 experiments were performed, with 14 Cells and
 309 3 Optérons operated in the beam. The first three experiments,
 310 which were the only data collected for Triblade 2, were omitted
 311 from the results since these tests had three Tribblades in the
 312 beam whereas the remaining experiments had only a single
 313 Triblade in the beam. Another experiment with missing beam
 314 fluence data was also omitted from the analyzes. The Opteron
 315 beampath tests were performed after the Cell beampath tests
 316 since the Cells were of primary interest in the testing. Thus, the
 317 behavior of the Optérons and the hardware in their beampaths

TABLE I
PROPORTION REDUCTION AT EACH OF FOUR BEAM AIMS

Upper Cell	Lower Cell	Upper Opteron	Lower Opteron
0.60	0.69	0.69	0.70
(0.56, 0.64)	(0.66, 0.72)	(0.56, 0.84)	(0.60, 0.79)

in a Triblade with no previous exposure to the beam cannot be
 estimated based on this testing.

Two different beam diameters were used for the experi-
 ments: a two-inch beam diameter for the first 53 experiments
 and a one-inch beam diameter for the remaining 59, where
 these beam diameter measurements reflect the full-width half-
 maximum (FWHM) boundary. All testing was performed at
 nominal voltages and nominal temperatures with the test fixture
 at normal incidence to the beam.

C. Corrected Neutron Fluences

Typically, corrected neutron fluences would be based on
 the decrease in flux given the distance from the beam source.
 Without the BC-H chassis and the Tribblades in the beam,
 the calculated decrease in flux from the beam source to the
 downbeam Virtex-II is about 20%.

The reduction in flux at each of the four beam aims de-
 scribed at the end of the Hardware Test Setup description in
 Section III.A (lower Cell, upper Cell, lower Opteron, upper
 Opteron) was estimated based on the experimental data to
 assess whether the BC-H and Triblade led to additional re-
 duction in beam intensity. Table I presents the posterior mean
 of the reduction in beam flux at each of the four beam aims
 and corresponding 95% credible intervals (CI). These values
 are based on the Virtex-II measurements from the upbeam
 and downbeam Virtex-II devices and a distance of 95 inches
 between the point at which the beam enters the test facility and
 the downbeam Virtex-II, which is the most common distance
 between these two points in the experimental data, and they
 incorporate both attenuation resulting from the material in the
 beam and divergence of the beam resulting from distance from
 the beam source. These results and those throughout this paper
 are based on the model described in Section IV-D, which
 permits different reductions in the beam at each beam aim and
 incorporates the uncertainty in these reductions; see [14] for
 details. The narrower CIs for the Cell beam aims reflect the
 greater number of experiments performed at the two Cell beam
 aims.

Since Table I demonstrates that the decrease in flux is larger
 than that expected due to only distance from the beam source,
 the neutron fluence exposures for different tests were corrected
 based on both distance from the beam source and attenuation
 through matter. The decrease in flux based on distance was
 calculated as usual, i.e., under the assumption that the beam
 is a point source with the reduction proportional to the squared
 distance from the beam source. The decrease in radiation due
 to attenuation through variable matter (i.e., the Triblade) is
 difficult, if not impossible, to account for precisely. However,
 as described in the previous paragraph the total proportion
 reduction through the entire Triblade and BC-H for each of the

TABLE II
HARDWARE NEUTRON EXPOSURE: AN ASTERISK (*) INDICATES A
QS22 THAT EXPERIENCED A VOLTAGE REGULATOR
FAILURE DURING THE TESTING

Blade	Beam Aim	Corrected Fluence ($\frac{neutrons}{cm^2}$)	
		Lower Bound	Upper Bound
Triblade 1 LS21	Upper Opteron	9.20×10^7	2.38×10^8
Triblade 1 LS21	Lower Opteron	3.42×10^8	8.94×10^8
Triblade 1 QS22a*	Upper Cell	1.28×10^9	2.56×10^9
Triblade 1 QS22a*	Lower Cell	7.98×10^8	2.08×10^9
Triblade 1 QS22b	Upper Cell	1.29×10^9	2.57×10^9
Triblade 1 QS22b	Lower Cell	8.01×10^8	2.09×10^9
Triblade 2 LS21	Upper Opteron	0	0
Triblade 2 LS21	Lower Opteron	0	0
Triblade 2 QS22a	Upper Cell	0	0
Triblade 2 QS22a	Lower Cell	8.16×10^8	2.13×10^9
Triblade 2 QS22b*	Upper Cell	0	0
Triblade 2 QS22b*	Lower Cell	8.18×10^8	2.13×10^9
Triblade 3 LS21	Upper Opteron	0	0
Triblade 3 LS21	Lower Opteron	0	0
Triblade 3 QS22a	Upper Cell	9.98×10^9	1.99×10^{10}
Triblade 3 QS22a	Lower Cell	1.86×10^9	4.85×10^9
Triblade 3 QS22b*	Upper Cell	1.00×10^{10}	1.99×10^{10}
Triblade 3 QS22b*	Lower Cell	1.86×10^9	4.86×10^9
Triblade 4 LS21	Upper Opteron	0	0
Triblade 4 LS21	Lower Opteron	4.37×10^8	1.14×10^9
Triblade 4 QS22a	Upper Cell	3.82×10^9	7.62×10^9
Triblade 4 QS22a	Lower Cell	1.43×10^{10}	3.73×10^{10}
Triblade 4 QS22b*	Upper Cell	3.83×10^9	7.64×10^9
Triblade 4 QS22b*	Lower Cell	1.43×10^{10}	3.74×10^{10}

four beam aims can be estimated with the Virtex-II readings. With this information, the neutron fluence to which a particular component is exposed prior to a particular error or operator decision to stop a test is assumed to lie between a lower bound and an upper bound explained below. The resulting uncertainty in neutron exposure of the component is explicitly incorporated in the model described in Section IV-D and reflected in the results presented here. The lower bound assumes that all of the attenuation caused by the beam passing through the BC-H and Triblade under test happened upstream of the component under test, while the upper bound assumes all attenuation happened downstream of the component under test. While neither of these cases may reflect the actual reduction due to attenuation of the beam, they best capture the knowledge of beam attenuation that resides in the experimental data without making any further assumptions. See [14] for details of the model used.

Table II details the posterior means of the upper bounds and lower bounds of the corrected neutron fluence, for neutrons with energies greater than 10 MeV, accumulated at each beam aim during the testing. The corrected fluences are based on the posterior mean estimate, which averages over the uncertainty in the attenuation for each beam aim. The data for each test or experiment are provided in Table V in the Appendix. As it was

TABLE III
ESTIMATED FAILURE FIT AND SDC FIT FOR CELL BEAMPATHS
AND OPTERON BEAMPATHS

	Failure FIT	SDC FIT
Cell Beampaths	172	7.2
95% CI	(92, 524)	(2.1, 26)
Opteron Beampaths	940	119
95% CI	(306, 2934)	(30, 453)

not possible to test one Cell in a Triblade without exposing the 390 second Cell in its beampath, the fluences include the exposure 391 gained when a Cell was running the idle test code while the 392 other Cell in its beampath was under test. 393

IV. RESULTS

A. Longevity of Hardware in the Beam, Post-Beam Testing and Root Cause Analysis of Permanently Failed Hardware

Some hardware experienced permanent failures relatively quickly upon exposure to the beam, while other hardware had greater longevity in the beam (see Table III). For example, QS22b on Triblade 2 was unable to boot after exposure to a corrected neutron fluence of *at most* $2.13 \times 10^9 \text{ neutrons/cm}^2$, while the lower Cell on QS22a on Triblade 4 remained operational after exposure to a corrected neutron fluence of *at least* $1.43 \times 10^{10} \text{ neutrons/cm}^2$. That said, Triblade 4 was tested with the one-inch beam diameter so it had less hardware in the beam than Triblade 2, which was tested with the two-inch beam diameter.

Following the beam testing, Triblades 1, 3 and 4 were tested in a production platform at LANL. (Triblade 2 had suffered damage through handling, and post-irradiation testing at LANL was not possible.) This testing used all of the applications from the beam testing, with the exception of the bottom Opteron in Triblade 4, which was not tested with the Opteron-only correlator code. Triblades 1, 3, and 4 each had a QS22 that would not boot. In addition, the QS22 in Triblade 1 that would boot could not communicate with the relevant Opteron.

After this post-irradiation testing was completed, all four Triblades were returned to IBM for root cause failure analysis. It was found that each of the 4 Triblades had a QS22 that had permanently failed. Further, these permanent failures were the result of voltage shorts in voltage regulators. Voltage regulators have been experimentally shown to experience single-event burnout (SEB) and single-event gate rupture (SEGR), both of which are destructive effects that can cause the system to be fully or partially unbiased, when exposed to thermal and fast neutrons [34]. The failed voltage regulators should not have been within the FWHM boundary of the beam unless the beam was mistargeted, so the neutron exposure they received should be less than that reported for the corresponding beam aims in Table II.

B. Failure Data

Each experiment was categorized as having one of two end states: 1) survival, meaning that the experiment ended when

the experimenter believed the application was still running or 2) failure, indicating that the application was no longer running at the end of the experiment, e.g., because of an application or a system crash. Since the output from the test applications appeared roughly every minute, it is possible that in some cases in which the system is deemed to have survived the experiment it had actually failed, but that failure was not detected before the experiment ended. Post-irradiation analysis showed that 79 of the 94 tests conducted on the Cells ended in failure, while all 14 tests conducted on the Opteron ended in failure.

Observed failures include application hangs, blades that spontaneously rebooted, and blades that became non-responsive. Investigation of the log data did not yield definitive root causes. Our hypothesis is that in most cases the hardware failed so completely and so quickly that no useful diagnostic information could be obtained.

C. Silent Data Corruption

In order to check for SDC, the computational test codes included a step in which the calculated answer was compared to the correct answer. Four SDCs were observed. Two SDCs occurred when a Cell was running a computational test code (one with VPIC and one with correlator) and two SDCs occurred when an Opteron was running the Opteron-only correlator test code. For the Cell beampaths, the median posterior probability that an error is an SDC rather than a failure (e.g., application or system crash) is 0.038 with 95% CI (0.011, 0.088), while for the Opteron beampaths it is 0.114 with 95% CI (0.035, 0.250). These estimates along with their corresponding uncertainty statements were obtained using standard Bayesian statistical methods for analyzing Binomial data [14]. The Opteron CI is wider because fewer experiments were conducted for the Opteron beampaths.

D. Failure FITs and SDC FITs for Cell and Opteron Beampaths

A statistical model that incorporated the upper and lower bounds on fluence until error (failure or SDC) and that accounted for the application used for each test, the Triblade under test, the beam aim (Cell beampaths or Opteron beampaths), and the beam diameter was fit to the experimental data [14]. All results presented below derive from this model and pertain to the conditions under which the experiments were conducted, with results likely to be obtained under other conditions less clear. Further, the results presented here are based on the experimental data collected at LANSCE and not on failures or SDCs observed in the Roadrunner platform. All results have been estimated via Markov Chain Monte Carlo [35].

Based on this modeling, Table II presents estimated failure FITs and SDC FITs for the Cell beampaths and the Opteron beampaths, along with 95% CIs that capture the uncertainty in the FIT estimates. These and all other FIT estimates presented in this work are based on one Cell idling while the other runs the VPIC test code (since it is most representative of a computational application that might be used in the field of those considered in our study) and the two-inch beam diameter.

TABLE IV
ESTIMATED FAILURE FIT AND SDC FIT FOR A TRIBLADE, A ROADRUNNER CU, AND ROADRUNNER

	Failure FIT	SDC FIT
Triblade	2.22×10^3	2.58×10^2
95% CI	$(8.06 \times 10^2, 7.06 \times 10^3)$	$(9.09 \times 10^1, 9.89 \times 10^2)$
Roadrunner CU	4.51×10^5	5.23×10^4
95% CI	$(2.28 \times 10^5, 1.19 \times 10^6)$	$(1.70 \times 10^4, 1.61 \times 10^5)$
Roadrunner	7.69×10^6	8.81×10^5
95% CI	$(3.86 \times 10^6, 1.90 \times 10^7)$	$(2.85 \times 10^5, 2.78 \times 10^6)$

They reflect the flux of neutrons in Los Alamos, NM that have energies greater than 10 MeV, which is estimated to be normal with mean $0.019 \text{ neutrons/cm}^2/\text{sec}$ [36] and standard deviation of $4.4\text{e-}4 \text{ neutrons/cm}^2/\text{sec}$ [3].

The FIT estimates and corresponding uncertainty intervals are calculated by standard Bayesian analysis techniques. Specifically, a Monte Carlo procedure is used that repeatedly generates values of parameters from their posterior distribution based on the statistical model described, each time calculating FITs based on the generated parameter values. That is, the expected number of failures in 10^9 hours will be different for different parameter values. Thus, this procedure reflects the uncertainty in FIT due to the uncertainty in the unknown model parameters and the uncertainty in the amount of neutron exposure to the hardware under test. The FIT estimate is taken to be the median of the FITs calculated from the generated parameter values, while the 0.025 and 0.975 quantiles are used for the bounds of the 95% CIs. See [14] for details.

From the results in Table IV, the Cell beampaths are less susceptible to neutron-induced errors than the Opteron beampaths. Care must be taken in interpreting this result since these beampaths include hardware in addition to the microprocessor that was running applications during the testing. That is, the values in the Table II cannot be interpreted as reflecting only the Cell and Opteron microprocessors. In particular, a small amount of the Opteron memory was in the beam when the Cells were being tested, with more exposure resulting when using the two-inch beam diameter as opposed to the one-inch beam diameter. Similarly, a small amount of Cell memory was in the beam when testing one of the Opteron, but the Opteron in a particular Triblade were tested after the Cells in that Triblade were tested. Using the two-inch beam diameter versus the one-inch beam diameter does not significantly change the hazard rate or instantaneous error rate (see Section IV.E), suggesting that any resulting effects in the Opteron memory are not likely to be substantial.

Thus, while this study underscores that there is almost certainly a difference in neutron susceptibility between the hardware in the Opteron beampaths and the hardware in the Cell beampaths, identifying the source of this difference with certainty is not possible. Since all of the hardware in the beampaths of each of the processors was irradiated, it could reflect neutron interactions with this hardware rather than the processors themselves. Assuming that most if not all neutron effects occurred in the processors it could reflect their process technologies (the Cell is 65nm SOI and the Opteron is 90nm

534 SOI), transistor counts, caches sizes, numbers of susceptible
535 states, architectural vulnerability factors [37], [38], architec-
536 tures (the Cell architecture is somewhat simpler than that of the
537 Opteron) or some other cause.

538 *E. Effects of Application, Beam Aim, Beam Diameter, and* 539 *Triblade Under Test on the Error Rate*

540 Based on the results of the model described in Section IV-D,
541 the paragraphs below discuss the effects of increasing exposure
542 to the beam, beam aim, Triblade under test, application used
543 for the test, and beam diameter on the hazard rate, i.e., the
544 instantaneous error (failure and SDC) rate of the hardware
545 under test.

546 The baseline hazard rate appears to be close to constant,
547 suggesting that the instantaneous error rate likely does not vary
548 much with increasing exposure to the beam for the exposures
549 observed in this study. Therefore, it is likely that sensitivity
550 to neutrons does not change with increasing dose accumula-
551 tion and in-field usage should have roughly constant neutron-
552 induced error rates.

553 The posterior probability that the beam aim (Cell beampaths
554 or Opteron beampaths) affects the hazard rate is 1.0, meaning
555 that there is most certainly a difference in neutron sensitivity
556 between the hardware in the Cell beampaths and the hardware
557 in the Opteron beampaths. With the Opteron beampaths, the
558 median multiplier to the hazard rate is 5.884 with 95% CI
559 (2.749, 11.753), meaning that errors are roughly six times
560 more frequent with the Opteron beampaths than with the Cell
561 beampaths.

562 There is a relationship between the Triblade under test
563 and the beam diameter used for the testing. Triblade 3 was
564 tested using the two-inch beam diameter and Triblade 4 was
565 tested using the one-inch beam diameter, while Triblade 1
566 was tested using both beam diameters. With a situation like
567 this, it can be difficult to determine which of Triblade under
568 test or beam diameter is more influential on the hazard rate.
569 That said, the posterior probability that one or both of Triblade
570 under test and beam diameter affects the hazard rate is 0.931,
571 and the results below suggest that Triblade under test is more
572 likely than beam diameter to affect the hazard rate.

573 The modeling results indicate a 0.897 posterior probability
574 that different Triblades under test experienced different sensi-
575 tivities to the beam. The posterior median relative difference
576 in hazard rate for two randomly-selected Triblades is 1.357
577 with 95% CI (1.000, 5.049). Thus, this test data suggests
578 that process-variation-based differences in neutron sensitivity
579 may exist. However, more Triblades would need to be tested
580 and/or more time would need to be spent under test to fully
581 investigate the implications of process-variation-based neutron
582 sensitivities.

583 Beam diameter (one-inch versus two-inch) has a 0.198 pos-
584 terior probability of affecting the hazard rate, suggesting that
585 beam diameter did not have much if any impact on the hazard
586 rate. This implies that most of the sensitive hardware likely lies
587 within the one-inch beam diameter.

588 For the most part, the application being run did not affect
589 the hazard rate. The largest effect on the hazard rate is for

hybrid Linpack, with a 0.417 posterior probability of having
a hazard rate different from that of the idle condition. Its
median multiplicative effect on the hazard rate is 1.000, with
95% CI (1.000, 2.545). Therefore, the error sensitivity did not
have much application dependence. This result is consistent
with related findings in [5] and the confidence limits presented
in [4].

There are a number of possible explanations for this result.
First, the operating system, which executed in all tests whether
an application was executing or not, might be overshadowing
the effect of the application on the hardware sensitivity to
neutrons. In [39] results from [16] are used to indicate that the
proton cross-section for the Pentium II and MMX microproces-
sors was two to three orders of magnitude larger when tested
with Windows operating system than without. Since definitive
root causes for observed failures could not be determined, it
could be that enough failures resulted from OS tasks rather
than application tasks that it is not possible to distinguish large
differences among the applications. Second, the applications
chosen here may have similar neutron sensitivities, which other
applications might not share. Further study with more appli-
cations with different programming and computing patterns
would be useful. To better understand the extent to which
failures derive from OS tasks, the testing could be performed
with the applications running on the processors under test, but
without an OS.

566 *F. Projected Failure and SDC Rates for Roadrunner*

Roadrunner is composed of 17 connected units (CU), each
of which includes 180 Triblades that are used for computation.
The experimental results can further be used to estimate failure
FITs and SDC FITs and corresponding 95% CIs for a single
Triblade, for the 180 Triblades in a CU, and for all of the
Triblades in the Roadrunner platform (17 CUs); Table IV
provides these values.

These results do not reflect the neutron sensitivity of all of
the hardware in a Triblade, as they only include the hardware
in the Cell and Opteron beampaths. For the Triblade values
they assume that errors in the hardware in the different beam-
paths occur independently, while the CU values further assume
independence of errors in the Triblades within a CU and the
Roadrunner values assume independence of all Triblades within
Roadrunner. See Section IV-D for additional assumptions un-
derlying these FIT estimates.

Table IV indicates that for a Triblade, Roadrunner CU, and
Roadrunner the failure FIT estimate is roughly an order of
magnitude larger than the SDC FIT estimate. Roadrunner is
estimated to experience one cosmic-ray-neutron-induced fail-
ure roughly every 130 hours of operation and one cosmic-ray-
neutron-induced SDC roughly every 1100 hours of operation.

The effect of any SDCs on calculations performed on Road-
runner is likely to be small since the results of many cal-
culations are typically combined to produce a final result,
thus mitigating the effect of an SDC in any one of the un-
derlying calculations. Specifically, verification and validation
efforts involve parameter studies that enable errors bars to be
investigated and better understood, with a suite of calculations,

TABLE V
EXPERIMENTAL DATA

Record	Hardware Tested	Application	SDC	Fluence A	Fluence B
1	cell: 3b-low	varied	0	4.81×10^8	1.51×10^9
2	cell: 3b-low	varied	0	1.21×10^8	5.04×10^8
3	cell: 3b-low	varied	0	8.21×10^7	3.10×10^8
4	cell: 3b-low	corr	0	3.56×10^6	1.12×10^8
5	cell: 3b-low	corr	0	7.98×10^6	1.22×10^8
6	cell: 3b-low	varied	0	4.23×10^7	1.90×10^8
7	cell: 3b-upp	varied	0	1.37×10^8	3.59×10^8
8	cell: 3b-upp	varied	0	1.42×10^8	3.73×10^8
9	cell: 3b-upp	varied	0	6.71×10^7	3.94×10^8
10	cell: 3b-upp	varied	0	5.09×10^8	1.16×10^9
11	cell: 3b-upp	vpic	0	6.78×10^7	2.13×10^8
12	cell: 3b-upp	hpl	0	0.00	9.34×10^7
13	cell: 3b-upp	hpl	0	1.31×10^8	3.52×10^8
14	cell: 3b-upp	hpl	0	6.72×10^7	2.21×10^8
15	cell: 3b-upp	hpl	0	4.87×10^6	3.07×10^8
16	cell: 3b-upp	corr	0	4.70×10^7	1.77×10^8
17	cell: 3b-upp	corr	0	4.30×10^8	9.44×10^8
18	cell: 3b-upp	corr	0	1.19×10^8	3.30×10^8
19	cell: 3b-upp	corr	0	8.56×10^8	Inf
20	cell: 3b-upp	vpic	0	2.99×10^7	1.21×10^8
21	cell: 3b-upp	vpic	0	4.52×10^7	9.05×10^7
22	cell: 3b-upp	vpic	0	5.31×10^7	1.06×10^8
23	cell: 3b-upp	vpic	0	2.00×10^8	4.38×10^8
24	cell: 3b-upp	vpic	0	9.51×10^8	2.05×10^9
25	cell: 3b-upp	idle	0	3.49×10^8	7.48×10^8
26	cell: 3b-upp	int_add	0	2.19×10^8	4.99×10^8
27	cell: 3b-upp	cg	0	7.26×10^7	Inf
28	cell: 3b-upp	cg	0	9.66×10^7	3.86×10^8
29	cell: 3b-upp	corr	0	7.06×10^7	3.11×10^8
30	cell: 3b-upp	idle	0	6.69×10^8	Inf
31	cell: 3b-upp	int_add	0	6.49×10^8	Inf
32	cell: 3b-upp	vpic	0	1.27×10^8	3.74×10^8
33	cell: 3b-upp	vpic	1	2.14×10^7	6.65×10^7
34	cell: 3b-upp	vpic	0	2.20×10^8	Inf
35	cell: 3a-upp	hpl	0	3.04×10^8	8.02×10^8
36	cell: 3a-upp	hpl	0	3.56×10^7	2.73×10^8
37	cell: 3a-upp	hpl	0	0.00	1.32×10^8
38	cell: 3a-upp	hpl	0	1.87×10^7	1.54×10^8
39	cell: 3a-upp	hpl	0	1.58×10^8	Inf
40	cell: 3a-upp	int_add	0	4.50×10^8	1.11×10^9
41	cell: 3a-upp	int_add	0	1.98×10^8	4.54×10^8
42	cell: 3a-upp	int_add	0	9.42×10^7	Inf
43	cell: 3a-upp	hpl	0	2.64×10^7	1.51×10^8
44	cell: 3a-upp	hpl	0	1.35×10^8	3.40×10^8
45	cell: 3a-upp	hpl	0	3.23×10^7	8.94×10^7
46	cell: 3a-upp	hpl	0	4.06×10^8	9.22×10^8
47	cell: 1a-upp	corr	0	9.18×10^6	8.51×10^7

TABLE V
(Continued). EXPERIMENTAL DATA

Record	Hardware Tested	Application	SDC	Fluence A	Fluence B
48	cell: 1a-upp	corr	0	2.91×10^8	6.63×10^8
49	cell: 1a-upp	corr	1	1.31×10^8	2.95×10^8
50	cell: 1a-upp	corr	0	1.55×10^8	7.63×10^8
51	cell: 1a-upp	cg	0	2.67×10^8	5.95×10^8
52	cell: 1a-upp	cg	0	4.84×10^7	1.92×10^8
53	opt: 1-top	corr	0	4.35×10^7	1.26×10^8
54	opt: 1-top	corr	1	4.92×10^6	4.26×10^7
55	opt: 1-top	corr	0	0.00	5.01×10^7
56	opt: 1-top	corr	0	1.53×10^6	4.95×10^7
57	cell: 4b-upp	corr	0	4.68×10^8	9.80×10^8
58	cell: 4b-upp	corr	0	5.83×10^8	Inf
59	cell: 4a-upp	cg	0	4.56×10^8	Inf
60	cell: 4a-upp	corr	0	2.69×10^8	6.16×10^8
61	cell: 4a-upp	corr	0	2.16×10^8	Inf
62	cell: 4a-upp	int_add	0	2.60×10^8	5.60×10^8
63	cell: 4a-upp	hpl	0	8.29×10^8	Inf
64	cell: 4a-upp	idle	0	2.56×10^7	2.00×10^8
65	cell: 4a-upp	idle	0	2.00×10^8	4.68×10^8
66	cell: 4a-upp	idle	0	4.14×10^7	1.73×10^8
67	cell: 4a-upp	cg	0	1.31×10^8	3.49×10^8
68	cell: 4a-upp	cg	0	3.34×10^7	1.23×10^8
69	cell: 4b-low	vpic	0	7.15×10^8	Inf
70	cell: 4b-low	cg	0	6.52×10^8	1.83×10^9
71	cell: 4b-low	corr	0	2.03×10^8	1.29×10^9
72	cell: 4b-low	corr	0	1.74×10^8	6.54×10^8
73	cell: 4b-low	corr	0	1.75×10^8	5.21×10^8
74	cell: 4b-low	int_add	0	5.36×10^8	1.47×10^9
75	cell: 4b-low	int_add	0	7.15×10^8	1.93×10^9
76	cell: 4b-low	hpl	0	4.99×10^7	1.75×10^8
77	cell: 4b-low	hpl	0	5.24×10^8	1.43×10^9
78	cell: 4b-low	idle	0	3.56×10^8	9.83×10^8
79	cell: 4b-low	vpic	0	7.94×10^8	2.17×10^9
80	cell: 4b-low	corr	0	4.24×10^7	4.67×10^8
81	cell: 4b-low	corr	0	1.82×10^8	6.01×10^8
82	cell: 4b-low	corr	0	2.20×10^8	6.62×10^8
83	cell: 4b-low	cg	0	6.51×10^8	2.07×10^9
84	cell: 4b-low	int_add	0	5.01×10^8	1.37×10^9
85	cell: 4b-low	hpl	0	4.33×10^7	1.99×10^8
86	cell: 4b-low	hpl	0	8.90×10^6	5.87×10^7
87	cell: 4b-low	hpl	0	5.97×10^7	2.33×10^8
88	cell: 4b-low	hpl	0	4.05×10^8	Inf
89	cell: 4b-low	vpic	0	2.61×10^8	Inf
90	cell: 4b-low	vpic	0	1.66×10^8	6.69×10^8
91	cell: 4b-low	vpic	0	3.56×10^6	1.08×10^8
92	cell: 4b-low	vpic	0	4.38×10^8	1.25×10^9
93	cell: 4b-low	int_add	0	8.48×10^8	Inf
94	cell: 4b-low	corr	0	1.87×10^9	Inf

TABLE V
(Continued). EXPERIMENTAL DATA

Record	Hardware Tested	Application	SDC	Fluence A	Fluence B
95	cell: 4b-low	cg	0	7.31×10^8	1.96×10^9
96	cell: 4b-low	cg	0	3.52×10^7	1.13×10^8
97	cell: 4b-low	cg	0	1.03×10^8	3.00×10^8
98	cell: 4b-low	cg	0	1.24×10^8	4.05×10^8
99	cell: 4b-low	cg	0	1.71×10^8	4.89×10^8
100	cell: 4b-low	cg	0	1.33×10^8	4.02×10^8
101	cell: 4b-low	cg	0	8.70×10^8	2.75×10^9
102	opt: 4-low	corr	0	1.72×10^7	1.58×10^8
103	opt: 4-low	idle	0	1.41×10^8	3.70×10^8
104	opt: 4-low	idle	0	6.00×10^7	1.57×10^8
105	opt: 4-low	corr	0	0.00	1.98×10^8
106	opt: 4-low	corr	0	3.80×10^6	1.59×10^8
107	opt: 4-low	idle	0	3.84×10^7	1.01×10^8
108	opt: 1-low	corr	0	7.04×10^6	6.12×10^7
109	opt: 1-low	idle	0	3.42×10^6	4.05×10^7
110	opt: 1-low	idle	0	3.24×10^7	1.75×10^8
111	opt: 1-low	corr	1	4.94×10^7	1.62×10^8
112	opt: 1-low	corr	0	3.67×10^7	1.75×10^8
113	opt: 1-low	corr	0	1.88×10^7	3.13×10^8

including some used to investigate error bars, used for decision making.

V. CONCLUSION

Replicates of two microprocessors, the IBM PowerXCell 8i and the AMD Opteron 2210 HE, along with the hardware in their respective beampaths, were tested at LANSCE for neutron sensitivities. These tests indicated that both microprocessor beampaths were susceptible to neutron-induced errors and that the Opteron beampaths were more sensitive to neutrons than the Cell beampaths as evidenced by the failure FIT and SDC FIT estimated for each of these beampaths. The data further provided some evidence for process-variation-based neutron sensitivity differences. Little application-based neutron sensitivity differences were found, with hybrid Linpack most likely to lead to a somewhat elevated hazard rate. The results suggest that failures, e.g., application and system crashes, occur roughly an order of magnitude more often than SDCs for the Triblades under test and for the Roadrunner platform that leverages them for computation.

APPENDIX EXPERIMENTAL DATA

Table V provides the experimental conditions pertaining to and data collected for each of the 113 errors analyzed for the results presented here. These errors include 109 experiments that ended with a failure or an operator decision to terminate the experiment and 4 SDCs. The data includes the following columns: Record (which corresponds to the sequential order in

which errors were observed and of tests that an operator ended while the system remained operational); Hardware Tested (the Triblade and location on that Triblade at which the beam was aimed; in the case of Cells running a computational code, it also provides which Cell was running the computational code, i.e., cell: 3a-upp, means Triblade 3 was in the beam, with the beam aimed at the upper Cells with the upper Cell in QS22a running a computational application); Application (the test code that was run prior to the error (crash or SDC) denoted as follows: hpl (hybrid Linpack), corr (correlator), cg (conjugate gradient), vplic (VPIC), integer; adder (int_add), varied (varied), and idle (idle)); SDC (a value of 1 indicates that an SDC occurred, with a 0 if otherwise), Fluence A (posterior mean of the lower bound for the neutron fluence for neutrons with energies above 10 MeV accumulated at the processor under test until error), and Fluence B (posterior mean of the upper bound for the neutron fluence for neutrons with energies above 10 MeV accumulated at the processor under test until error, with a value of "Inf" indicating that the operator decided to terminate the experiment prior to an error occurring).

ACKNOWLEDGMENT

The authors would like to thank J. Abeyta, C. Alexander, B. Bergen, A. Borrett, H. Brandt, J. Campa, R. Cardon, N. De Bardeleben, T. Fairbanks, P. Fields, A. Gibson, G. Grider, J. Loncaric, P. Lujan, A. Malin, F. Marshall, A. Montoya, J. Morrison, A. Shewmaker, M. Vigil, B. Villa, S. Wender, A. White, and C. Wright. The authors apologize for any inadvertent omissions from this list. The authors further thank the reviewers for their insightful comments.

The U.S. Government has rights to use, reproduce, and distribute this information. The public may copy and use this information without charge, provided that this Notice and any statement of authorship are reproduced on all copies. Neither the Government nor LANS makes any warranty, express or implied, or assumes any liability or responsibility for the use of this information. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness. This paper is published under LA-UR-12-20226. Copyright copyright 2010 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained by sending a request to pubpermissions@ieee.org.

REFERENCES

- [1] R. Baumann, "Radiation induced soft errors in advanced semiconductor technologies," *IEEE Trans. Device Mater. Rel.*, vol. 5, no. 3, pp. 305–316, Sep. 2005.
- [2] C. Constantinescu, "Neutron SER characterization of microprocessors," in *Proc. Int. Conf. Dependable Syst. Netw.*, 2005, pp. 754–759.
- [3] S. Michalak, K. Harris, N. Hengartner, B. Takala, and S. Wender, "Predicting the number of fatal soft errors in Los Alamos National Laboratory's ASC Q supercomputer," *IEEE Trans. Device Mater. Rel.*, vol. 5, no. 3, pp. 329–335, Sep. 2005.
- [4] H. Ando, R. Kan, Y. Tosaka, K. Takahisa, and K. Hatanaka, "Validation of hardware error recovery mechanisms for the SPARC64 V microprocessor," in *Proc. IEEE Int. Conf. Dependable Syst. Netw.*, 2008, pp. 62–69.

- [5] S. Rao, T. Hong, P. Sanda, J. Ackaret, A. Barrera, J. Yanez, S. Mitra, J. Kellington, and R. McBeth, "Examining workload dependence of soft error rates," in *Proc. SELSE*, 2008.
 - [6] P. Sanda, J. Kellington, P. Kudva, R. Kalla, R. McBeth, J. Ackaret, R. Lockwood, J. Schumann, and C. Jones, "Soft_error resilience of the IBM POWER6 processor," *IBM J. Res. Develop.*, vol. 52, no. 3, pp. 275–284, 2008.
 - [7] T. Hong, S. Michalak, T. Graves, J. Ackaret, and S. Rao, "Neutron beam irradiation study of workload dependence of SER in a microprocessor," in *Proc. SELSE*, 2009, pp. 1–5.
 - [8] F. Wrobel, J. Gasiot, and F. Saigne, "Hafnium and uranium contributions to soft error rate at ground level," *IEEE Trans. Nucl. Sci.*, vol. 55, no. 6, pp. 3141–3145, Dec. 2008.
 - [9] J. Autran, D. Munteanu, P. Roche, G. Gasiot, S. Martinie, S. Uznanski, S. Sauze, S. Semikh, E. Yakushev, S. Rozov, P. Loaiza, G. Warot, and M. Zampaolo, "Soft_errors induced by terrestrial neutrons and natural alpha_particle emitters in advanced memory circuits at ground level," *Microelectron. Reliab.*, vol. 50, no. 9_11, pp. 1822–1831, 2010.
 - [10] Z. Kalbarczyk, R. Iyer, and L. Wang, "Application fault tolerance with armor middleware," *IEEE Internet Comput.*, vol. 9, no. 2, pp. 28–37, Mar./Apr. 2005.
 - [11] D. Katz, *Application-Based Fault Tolerance for Spaceborne Applications*. [Online]. Available: <http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/10574/1/02-2631.pdf>
 - [12] K. Koch, *Roadrunner Platform Overview*, 2008. [Online]. Available: <http://www.lanl.gov/orgs/hpc/roadrunner/pdfs/Koch%20-%20Roadrunner%20Overview/RR%20Seminar%20-%20System%20Overview.pdf>
 - [13] H. Meuer, *31st TOP500 List Topped by First_Ever Petaflop/s Supercomputer Scientific Computing*, 2008. [Online]. Available: <http://www.scientificcomputing.com/31st-TOP500-List-Topped-by-First-ever-Petaflops-Supercomputer.aspx>
 - [14] C. Storlie, S. Michalak, H. Quinn, A. DuBois, and D. DuBois, "A Bayesian survival analysis of neutron-induced errors in high performance computing hardware," Los Alamos Nat. Lab., Los Alamos, NM, 2011.
 - [15] S. Michalak, A. DuBois, C. Storlie, H. Quinn, W. Rust, D. DuBois, D. Modl, A. Manuzzato, and S. Blanchard, "Neutron beam testing of high performance computing hardware," in *Proc. IEEE Nucl. Space Radiation Effects Conf. Radiation Effects Data Workshop*, 2011, pp. 1–8.
 - [16] D. Hiemstra, A. Baril, and M. Dettwiler, "Single event upset characterization of the Pentium (R) MMX and Pentium (R) II microprocessors using proton irradiation," *IEEE Trans. Nucl. Sci.*, vol. 46, no. 6, pp. 1453–1460, Dec. 1999.
 - [17] D. Hiemstra and A. Baril, "Single event upset characterization of the Pentium (R) MMX and Celeron (R) microprocessors using proton irradiation," in *Proc. IEEE Radiation Effects Data Workshop*, 2000, pp. 39–44.
 - [18] D. Hiemstra, S. Yu, and M. Pop, "Single event upset characterization of the Pentium (R) MMX and Low Power Pentium (R) MMX microprocessors using proton irradiation," in *Proc. IEEE Radiation Effects Data Workshop*, 2001, pp. 32–37.
 - [19] D. Hiemstra, S. Yu, and M. Pop, "Single event upset characterization of the Pentium 4, Pentium III and Low Power Pentium MMX microprocessors using proton irradiation," in *Proc. IEEE Radiation Effects Data Workshop*, 2002, pp. 51–57.
 - [20] G. Swift, F. Farmanesh, F. Guertin, F. Irom, and D. Millward, "Single_event upset in the PowerPC750 microprocessor," *IEEE Trans. Nucl. Sci.*, vol. 48, no. 6, pp. 1822–1827, Dec. 2001.
 - [21] F. Irom, F. Farmanesh, and C. Kouba, "Single_event upset and scaling trends in new generation of the commercial SOI PowerPC microprocessors," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 6, pp. 3563–3568, Dec. 2006.
 - [22] F. Irom, T. Miyahira, D. Nguyen, I. Jun, and E. Normand, "Results of recent 14 MeV neutron single event effects measurements conducted by the Jet Propulsion Laboratory," in *Proc. IEEE Radiation Effects Data Workshop*, 2007, pp. 141–145.
 - [23] A. Dixit, R. Heald, and A. Wood, "Trends from ten years of soft error experimentation," in *Proc. SELSE*, 2009, pp. 1–4.
 - [24] A. Dixit and A. Wood, "The impact of new technology on soft error rates," in *Proc. SELSE*, 2011, pp. 486–492.
 - [25] M. Gschwind, P. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, and T. Yamazaki, *A Novel SIMD Architecture for the Cell Heterogeneous Chip Multiprocessor*, 2005. [Online]. Available: http://www.hotchips.org/archives/hc17/2_Mon/HC17.S1/HC17.S1T1.pdf
 - [26] AMD, *The AMD Opteron Processor for High-End Embedded Designs*, 2008. [Online]. Available: http://www.amd.com/es/Documents/42350B_Operon_PB.pdf
 - [27] IBM, *Bladecenter H Chassis*. [Online]. Available: <http://www-03.ibm.com/systems/bladecenter/hardware/chassis/bladeh/>
 - [28] *Virtex-II Platform FPGAs : Complete Data Sheet (DS031)*. [Online]. Available: <http://www.xilinx.com>
 - [29] H. Quinn, P. Graham, M. Wirthlin, B. Pratt, K. Morgan, M. Caffrey, and J. Krone, "A test methodology for determining space readiness of Xilinx SRAM_based FPGA devices and designs," *IEEE Trans. Instrum. Meas.*, vol. 58, no. 10, pp. 3380–3395, Oct. 2009.
 - [30] M. Kistler, J. Gunnels, D. Brokenshire, and B. Benton, "Programming the Linpack benchmark for Roadrunner," *IBM J. Res. Develop.*, vol. 53, no. 5, pp. 736–746, Sep. 2009.
 - [31] A. Petitet, R. C. Whaley, J. Dongarra, and A. Cleary, *HPL_A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers*. [Online]. Available: <http://netlib.org/benchmark/hpl/>
 - [32] A. DuBois, C. Conner, S. Michalak, G. Taylor, and D. DuBois, "Application of the IBM cell processor to real-time cross-correlation of a large antenna array radio telescope," Los Alamos Nat. Lab., Los Alamos, NM, Tech. Rep. LA-UR-09-03483, 2009.
 - [33] K. J. Bowers, B. J. Albright, L. Yin, W. Daughton, V. Roytershteyn, B. Bergen, and T. J. T. Kwan, "Advances in petascale kinetic plasma simulation with VPIC and Roadrunner," *J. Phys., Conf. Ser.*, vol. 180, no. 1, p. 012055, 2009.
 - [34] A. Hands, P. Morris, K. Ryden, C. Dyer, P. Truscott, A. Chugg, and S. Parker, "Single event effects in power MOSFETS due to atmospheric and thermal neutrons," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 6, pp. 2687–2694, Dec. 2011.
 - [35] A. Gelman, J. Carlin, H. Stern, and D. Rubin, *Bayesian Data Analysis*, 2nd ed. London, U.K.: Chapman & Hall, 2004.
 - [36] *Flux Calculation*, 2008. [Online]. Available: <http://seutest.com/cgi-bin/FluxCalculator.cgi/>
 - [37] S. Mukherjee, C. Weaver, J. Emer, S. Reinhardt, and T. Austin, "A systematic methodology to compute the architectural vulnerability factors for a high-performance microprocessor," in *Proc. 36th Annu. Int. Symp. MICRO*, 2003, pp. 29–40.
 - [38] A. Biswas, P. Racunas, R. Cheveresan, J. Emer, S. Mukherjee, and R. Rangan, "Computing architectural vulnerability factors for address_based structures," in *Proc. 32nd ISCA*, 2005, pp. 532–543.
 - [39] F. Irom, "Guideline for ground radiation testing of microprocessors in the space radiation environment," Jet Propulsion Lab., Los Angeles, CA, JPL Pub. 08-13, 2008.
- Sarah E. Michalak**, photograph and biography not available at the time of publication.
- Andrew J. DuBois**, photograph and biography not available at the time of publication.
- Curtis B. Storlie**, photograph and biography not available at the time of publication.
- Heather M. Quinn**, photograph and biography not available at the time of publication.
- William N. Rust**, photograph and biography not available at the time of publication.
- David H. DuBois**, photograph and biography not available at the time of publication.
- David G. Modl**, photograph and biography not available at the time of publication.
- Andrea Manuzzato**, photograph and biography not available at the time of publication.
- Sean P. Blanchard**, photograph and biography not available at the time of publication.

AUTHOR QUERY

NO QUERY.